

USER INTERFACE, OPERATING SYSTEM AND ARCHITECTURE

Cross-Reference to Related Applications

5 The present application for patent incorporates by reference and claims the priority of commonly owned United States Provisional Applications for Patent Serial No. 60/409,483 filed September 10, 2002 and Serial No. 60/428,163 filed November 21, 2002.

Background of the Invention

10 In recent years, the mobile computing marketplace has been flooded with smaller and faster personal devices. These next generation handhelds, cell phones, automotive telematics and personal information devices sport high resolution screens in excess of 320 by 240 pixels, processors faster than 300 MHz, and more than 16 Mbyte of dynamic RAM. New cell phones already in the marketplace combine traditional
15 cellular technology with the power of a small personal computer. Future devices will create brand new breeds of handheld personal computers with unparalleled power, mobility and battery life.

 Despite these impending revolutions in device capability, the market has failed to produce common, underlying standards for mobile interfaces. Even in this nascent
20 stage of device adoption, it is clear to industry observers that the wireless application market has become fragmented. Java-based technologies such as J2ME have been stymied with high hurdles to application qualification and numerous manufacturer-specific additions. Microsoft .NET technology, while promising, has been slow to grow on non-Intel platforms. Binary-based solutions like SymbianOS and BREW have failed to
25 develop traction in a cautious market. Competing standards and interests in the industry among carriers, device manufactures, technology vendors, operating system vendors and application developer continue to plague the industry.

 Several significant hurdles obstruct easy development of rich, widely deployable applications that run (at least partially) natively on a device:

- Multiple platforms require porting code to a variety of different API standards. Even within a single architecture specification, different platforms have slight variations that require extensive testing and modification.
- Competing languages (primarily C and Java) have no easy porting relationship.
- 5 • Service operators have created qualification hurdles that require expensive certification tests for applications before they can be deployed on a given platform.
- The industry has yet to agree on a common distribution mechanism for applications. Consequently, getting the application onto multiple devices
10 becomes extremely difficult.

Various sets of interests must be satisfied before these devices will realize their widespread potential as business and consumer platforms: For example, end users seek a rich, easy-to-use interface and a cost effective and easy way to acquire and run
15 mobile applications. Furthermore, they seek a user interface that is common across applications. Users want some reasonable assurances about the quality and integrity of the applications they receive. They are also interested in making greater use of the power and flexibility of the new devices.

Carriers, mobile operators, media companies and other entities seeking to
20 provide and distribute applications must choose a means to encourage application developers while maintaining their own revenue streams. Also carriers maintain a high concern about the security of code running on the devices connected to their networks

Developers (whether independent or employed by software companies) need
25 platforms that are economically viable to develop upon, non-complicated and connected to distribution points or mechanisms capable of disseminating applications widely. Many developers are small software shops that cannot afford to develop on multiple platforms and seek multiple certifications from operators.

It would therefore be desirable to provide methods and structures that would address these concerns.

In this regard, it is worth noting that most computers in common use are based around the general computing model established by John Von Neumann in 1945. In this model, programs are written in the form of ordered lists of instructions and stored within the main memory of a computing system. Variations on this design since 1945 have led to the architecture of modern Intel, PowerPC and other microprocessor-based systems. Even the Java platform by Sun Microsystems, which uses a virtual machine that runs on host platforms, uses a variation on the Von Neumann architecture.

One implication of the Van Nueman machine is an opcode-based program execution model, in which instructions are written for the native processing language of the platform (or a virtual processing language in the case of Java). These instructions consist of programming opcodes that have varying degree of granularity. In some cases they can be as simple as very basic mathematics and memory loading and saving. In other cases, they are more advanced, where individual opcodes for object manipulation, memory allocation and I/O systems.

In this traditional execution model, a processor reads instructions from a memory location, decodes the instructions found there and performs some work, which may involve additional reading and writing of memory data. When the instruction is complete, and succeeding instruction is read, either the next instruction in sequence or an instruction at the beginning of a new sequence. In this way, complicated program behavior can be established through creation of algorithms, loops, conditional statements, etc. Those familiar with the arts of computer design are aware that modern implementations use extremely complex optimizations and heuristics to achieve greater performance in this basic model, including Intel's Superscalar architecture.

To ease the development of programming instructions for a computer application, programs called compilers translate a program from a high-level representation available in structured text, XML or some other human-significant format to the lower level machine language of the underlying hardware (or emulation of hardware.) In this case, a programming language with certain grammatical rules is established, and programs are written to conform to this specification. Examples of compiled languages include C, C++, Java, Fortran, and Cobol.

In the traditional execution model, programs wishing to access resources outside of their memory space must use a set of conventions supported by an operating system. These conventions allow each program that runs on a computer to access resources in a consistent and protectable manner. Such a convention may be regarded
5 as an Application Programmer Interface (or API.) Typically, the API calls are triggered either by joining a program with libraries from the operating system at link-time or at run-time, or through the low-level triggering of interrupts, processor exceptions, or special memory regions that contain significance to the host operating system. Over time, these conventions have become quite extensive, where the typical operating
10 system (such as Windows 2000) contains tens of thousands of application programmer interface hooks, accessible either by shared library or through compile-time insertion of code. These APIs are typically an arbitrary organization of computer and operating system functionality and contain special rules. For instance, `setWindowTitle()` in Windows cannot be called before a `newWindow()` call has been made. The programmer
15 is responsible for being consistent both to the rules of the language as well as the rules of the API.

In some cases, programs written at a high level are not compiled down into their machine-language equivalent. Instead, a native computer program written in machine-language acts as a emulating processor, reading an internal representation of a current
20 instruction, executing it and then following its own rules for the selection of a subsequent instruction. These interpreted programming languages offer the benefit of greater flexibility for the language designer, easier implementation of a new programming language, usually at the expense of some efficiency since the a level of translation is always in place between the hardware processor and the internal
25 instruction processor. Examples include Perl, versions of Basic, UNIX shell languages, Python, etc.

Typical computer applications organize their information within their memory space using arbitrary data structures, which divide that space into meaningful chunks of data that can be referenced implicitly by the loading and saving operations of the
30 processor. To the computer and to the operating system, these data structures are

entirely private, and only the continuing flow of instructions through the processor gives these organizations any meaning. If a program wishes to save its context and state to resume execution later, it must conform to one of two constraints.

5 It must marshal up all of its relevant data and innermost state pertinent to the resumption of the program in a meaningful way, into a private or publicly structured serial bundle of data and save that data to a persistent location. When the program must resume, it must obtain information to resurrect its internal data structures in their previous state by decomposing the bundle of data it previously wrote. Variations on this model include placing data into a database, serial memory ring or other external
10 structure. This strategy has the advantage that the state of the program could be resumed on a different platform or computer system than the original, however the application must take excessive care to ensure that every piece of data has been saved in an identifiable format.

15 It must save its existing buffer of memory, exactly as it was left, into a single large snapshot, disentangle any memory references maintained by the operating system and write the complete state of the processor to disk. A variation on this method is used by Microsoft's operating system in its hibernation mode, and by the old BSD tool "undump." This method requires very little work to decompose and marshal private data structure into a secondary form, however it contains a very stringent
20 limitation: The application "image" can only be restored on a system that implements the same hardware and operating system state that the original program had run on. Any differences could cause unpredictable problems with the executable.

These complexities express themselves in the equally complex strategies computer programmers have used to construct computer programs that run on multiple
25 systems simultaneously. For example, when a small, embedded device needs to perform some complex processing operation on a more powerful server, it must build a complete description of the work required, and then transmit it to the remote server. This involves yet another API between the small device and the remote device (e.g., SOAP, and the like).

Thus, the Von Neumann model and the program execution models that result from it, while useful and ubiquitous, present limitations that computer scientists and engineers have recognized but failed to address. Some of these limitations include:

- 5 • **Instruction Based Programming:** Applications written for instantiations of the Von Neumann architecture must write their applications in the native (or emulated) code set of the given platform. Consequently, the ability for optimization after the code has been compiled lies within local heuristics usually available only to the instruction processor. Furthermore, the instructions must be
10 processed linearly, which encourages a procedural style of programming where program code is largely instructive instead of descriptive.
- **Seeding Complex Data Structures:** Applications that work with and manipulate complex nested data structures must populate those structures with either long lists of instructions (such as the list of instructions used to populate a
15 user interface,) or the loading of those data structures from some other source. Vast modifications to data structures after they have been initialized requires algorithms to be created within the program to manage the manipulations of the data. These algorithms can be complex and rely on an instructive model whereby the changes are made linearly step-by-step.
- 20 • **Remote Execution:** Remote execution involves either migrating images of a running program to a second location, or through the explicit marshalling and unmarshalling of relevant computing data
- **State Saving:** Programs that are halted and wish to save their state for resumption at a later time must use one of the two difficult strategies listed
25 above, each with potential drawbacks.
- **Private Application Structures:** Applications store all of their data using implicit data structures within their memory core. Consequently, applications that wish to share data with each other or with the host operating system must perform a translation into some common format (such as published data
30 structures, or XML, or whatever.)

- **API Issues:** Each operating system and software component must provide an application programmer interface to bridge data and activities between the program and resources available to the computer, including the user interface.

5 Thus, it would be desirable to provide methods and structures that overcome these limitations as well.

Summary of the Invention

10 The present invention addresses these limitations of the prior art, by providing methods and structures, aspects of which will be referred to variously hereinafter as "the Simple Quick Go system", "SQGO," "SimpleOperatingSystem", "SimpleOS" and the like, that enable simple, cost-effective design, construction, deployment and distribution of rich applications for computing devices, particularly suited for wireless devices and other mobile computing devices and platforms.

15 One aspect of the invention provides a method for creating and managing platform-independent applications. In one practice of the invention, this method includes: generating a platform-independent data superstructure defining the appearance and behavior of an application, independent of the characteristics of a digital processing device on which the application is to be instantiated; instantiating the
20 superstructure in the device, thereby installing the application for execution in accordance with the superstructure; updating, in response to events (which may include events generated by the application (indicative of application state), or by user or network input), information in a segment of the superstructure; and updating, in accordance with the superstructure segment update, the application state in the device.

25 In one aspect of the invention, the superstructure is an XML information structure. Application appearance and behavior are encapsulated within the superstructure, and application events are expressed to the superstructure via a pathway including a device-native operating system (OS) and a superstructure-dedicated OS acting as an intermediary between the device-native OS and the
30 superstructure. In this way, a defined portion of the application can be addressed and

updated in response to application events without necessitating update of the entire application, and the appearance and behavior of the application can be propagated with consistency across heterogeneous device types, to enable cross-device interoperability, replicability, and compatibility of applications and data with a consistency of user
5 experience.

Among other aspects that will be discussed in detail below, the invention also includes a programming language adapted for effectively constructing, interacting with and manipulating the superstructure and the SimpleOS.

In another aspect, the application instantiated in the device includes an
10 associated user interface whose behavior and state are defined by the superstructure. In a wireless or other handheld device, the interface can be advantageously adapted to be manipulated by a simple four-way controller actuated by a user's thumbs.

In a further aspect, methods are provided for device-to-device or cross-network communications and application deployment using the superstructure described in
15 detail below.

It should be noted that SimpleOS is but one implementation of the broader concepts of the invention. An operating system with a different design could also be used to implement the broad concepts described herein, particularly that of a Superstructure-Based Application Environment ("SBAE"). In other words, the structures
20 and methods described and illustrated by way of example herein are but selected implementations of the broad concepts of the invention. Many variations are possible and within the spirit and scope of the invention. Further aspects and implementation examples of the invention will next be set forth, in connection with the various drawing figures.

Brief Description of the Drawings

The foregoing summary and the following detailed description of the illustrated embodiments are better understood when read in conjunction with the appended
5 drawings, in which:

FIGS. 1A-3 are flowcharts depicting methods in accordance with the invention.

FIG. 4 is a schematic diagram depicting a network utilizing the present invention.

FIG. 5 is a schematic diagram depicting a client device utilizing the invention.

10 FIG. 6 is a schematic diagram depicting a server that can be used in conjunction with the invention.

FIG. 7 is a depicting of an example of the SimpleOS on a client device.

FIG. 8 is a schematic diagram depicting SimpleOS Portability.

FIG. 9 is a schematic diagram of the superstructure utilized in an embodiment of the invention.

15 FIG. 10 is a schematic diagram showing the organization of an object.

FIG. 11 presents an overview of a typical SQUI screen.

FIG. 12 is a screenshot of an example static application.

FIG. 13 is a screenshot of the Hello World example.

FIG. 14 is a screenshot of expression interpolation.

20 FIG. 15 is a screenshot relating to objects with variables.

FIG. 16 is a screenshot relating to variable incrementing.

FIG. 17 is a screenshot for the pizza ordering example.

FIG. 18 is a screenshot for text field extraction.

FIG. 19 is a screenshot showing further detail of a pizza ordering example.

25 FIG. 20 illustrates metaphors in client computing.

FIG. 21 depicts examples of various client/computing models.

Brief Description of Tables

Table 1 depicts a stylesheet example.

Table 2 depicts a simple object with a data tree.

5 Table 3 shows examples of identifiers.

Table 4 depicts examples of object data members.

Table 5 depicts defining methods for an object.

Table 6 shows a sample definition of an SQCard.

Table 7 depicts a sample frame definition.

10 Table 8 shows an example of defining the initial card.

Table 9 presents a static SimpleOS application example in SQML.

Table 10 shows four core programming modes.

Table 11 shows phases of the execution cycle.

Table 12 provides a synopsis of event handler modes.

15 Table 13 depicts the availability of modes within code contexts.

Table 14 shows an example of in-line event registration with code.

Table 15 shows in-line event registration for a method.

Table 16 shows in-line event registration to a server.

Table 17 depicts fields in a message descriptor.

20 Table 18 provides an example of a basic message descriptor.

Table 19 provides an example of combining in-line registrations with message descriptors.

Table 20 provides further detail regarding the Hello World example.

Table 21 depicts SQML tag attributes.

25 Table 22 shows an example of expression interpolation in text elements.

Table 23 gives an example of expression interpolation in attributes.

Table 24 shows an example of an object with a single variable.

Table 25 provides an example of an object with a string variable.

Table 26 provides an example relating to an object with a numeric variable.

30 Table 27 relates to an example of variable incrementing.

Table 28 provides a summary of examples of globally available variable attributes.

Table 29 relates to scalar attributes.

Table 30 relates to list methods.

5 Table 31 provides an example relating to the use of "This".

Table 32 shows declaring list variables.

Table 33 depicts a pizza ordering example.

Table 34 shows building a list of objects.

Table 35 shows text field extraction.

10 Table 36 provides examples of mathematical expressions.

Table 37 provides an example of instance variables.

Table 38 shows method definition.

Table 39 provides an example of evaluation tags in a method.

Table 40 shows method invocation.

15 Table 41 provides method invocation examples.

Table 42 shows variable masking.

Table 43 shows variable masking with element identifiers.

Table 44 provides examples of declaration order.

Table 45 shows blocks and variable scope.

20 Table 46 provides an example of an IF block.

Table 47 shows further detail of the pizza ordering example.

Table 48 gives an example of the body of a message.

Table 49 provides a summary of intrinsic methods.

Table 50 shows an example of serialized data.

25 Table 51 shows using fragments to update an application.

Table 52 relates to nodes.

Table 53 provides a summary of security threats and responses.

Table 54 provides a simple stylesheet example.

Table 55 shows a complex stylesheet example.

30 Table 56 provides a comparison of remote computing paradigms.

Detailed Description of the Invention

The following detailed description of the invention is organized into the following sections:

- I: Overview
- II: SimpleOS Software Concepts
- III: Composition of the Superstructure
- IV: Programming SimpleOS Applications
- V: SQScript Language
- VI: Superstructure Operations
- VII: Network Communication and Remote Execution
- VIII: Security
- IX: Stylesheets
- X: Additional Features

I. Overview

In one aspect, the invention provides a method for enabling the creation and management of platform-independent applications. In one practice of the invention, the method (see FIG. 1A) can include generating a platform-independent data superstructure defining the appearance and behavior of an application independent of characteristics of a digital processing device on which the application is to be instantiated (101), instantiating the application-defining superstructure (and thereby the application) on a device in accordance with the superstructure so that it can then be activated or run (102); validating the superstructure (103); generating operations (i.e., running the application) in the device to carry out the behavior defined or described by the current superstructure, thereby generating events (104); providing (by way of the user or a network) additional events (105); receiving the events, updating the superstructure in accordance with the events, and executing a handler (local or remote) which may update the superstructure (106); repeating the cycle by returning to validation (103); and as appropriate, updating, in accordance with the superstructure update, the application state in the device (108). These aspects are described in greater detail below.

In one practice of the invention, the superstructure is a hierarchical information structure (see FIGS. 9 and 10, discussed in greater detail below). More particularly, the

superstructure can be an XML structure. Application appearance and behavior are encapsulated within the superstructure, and application events are expressed to the superstructure via a pathway including a device-native operating system (OS) and a superstructure-dedicated OS acting as an intermediary between the device-native OS and the superstructure (FIG. 1A, 106). In this way, a defined portion of the application can be addressed and updated in response to application events without necessitating update of the entire application, and the appearance and behavior of the application can be propagated with consistency across heterogeneous device types, to enable cross-device interoperability, replicability, and compatibility of applications and data with a consistency of user experience.

In a further aspect of the invention, the superstructure can be serialized in whole or in part at any time (101).

A method according to the invention can also include (see FIG. 1B) generating a message containing a data object of a defined type operable to instantiate the application in a device (110), transmitting the message to a device operable to instantiate the application in accordance with the data object (112), receiving the message at the device (114), and instantiating the application in the device in accordance with the data object in the received message (116). The instantiating of the superstructure inside the target device can occur substantially when the application is invoked, or at an application provisioning time prior to application run-time (116).

In a further aspect of the invention (FIG. 1C), a provisioning application on a first device can locate within its operating environment a first superstructure for a new application superstructure to be expressed to a second device (118). The provisioning application can then generate a defined data object to be used to express the new application superstructure to the second device (120). The data object is transmitted to the second device via a message (122), and the second device creates a new application superstructure from the data object in the message (124). (In accordance with the invention it is contemplated that the latter can be accomplished by any operating system, or by SimpleOS.)

Alternatively, the provisioning application on a first device can locate within its operating environment a predefined data object that expresses a new application superstructure for a second device (118). The predefined data object can then be transmitted to the second device via a message (120), and the second device creates
5 its own copy of the new application superstructure from the data object in the message (122, 124). (Here again, in accordance with the invention it is contemplated that this latter can be accomplished by any operating system, or by SimpleOS.)

In a further practice of the invention (FIG. 1D), a first device can maintain an application capable of accepting input from a user to create an interactive message
10 (126). The first application can then translate an operational portion of the message into a new superstructure-based application operable to display the message, or cause interactive operations within the message (128); and the first application then initiates the transmission of the superstructure of the new application to a receiving device (130). The transmission of the superstructure can include converting the
15 superstructure into a temporary form that is transmitted, received, and decoded back into an original form on the receiving device (130); and the receiving device can maintain an application that receives the superstructure in its temporary form, decodes it, and causes the message-bearing superstructure to operate, thereby rendering the message (132).

20 In addition, a first device can transmit to a second device a message containing an application event item (140), and thereby cause the second device to place the application event item into a processing queue of the second device (142). This can be, for example, the SimpleOS processing queue.

In describing the operation, function and interpretation of the superstructure, it
25 is useful to note a distinction between the superstructure itself (a hierarchical information structure with various properties) and the "grammar" of the superstructure (which, by convention, describes a mapping between the data in the superstructure, and the behavior of an application). That mapping between data and behavior can remain the same across implementations of SimpleOS to ensure that the

superstructure, and applications operating in accordance with the superstructure, retain a consistent look and feel.

5 In accordance with the invention, for a given state of a selected application, the organization of the superstructure can be substantially invariant, regardless of the device, platform or device-native operating system environment in which the associated application is instantiated, so as to maintain a consistent application appearance and behavior across heterogeneous devices, platforms or device-native operating system environments, without substantial changes in the code or data structures of the application. The design of the superstructure can define rules of appearance and
10 behavior of the application that are substantially invariant across heterogeneous devices, platforms or device-native operating system environments; and substantially identical application source code can be used across heterogeneous devices, platforms or device-native operating system environments. Using the superstructure concept, a user interface (or GUI) provided by the application can have a substantially identical
15 appearance and behavior across heterogeneous devices, platforms or device-native operating system environments. It is not simply that applications will have a consistent appearance, but that because the superstructure maintains a "clean" abstraction around a common user interface (assisted by stylesheets as described below), this mechanism can ensure that the same superstructure will result in identical appearance and
20 operation on any device.

In one practice of the invention, operation of the application is implemented through operations on the superstructure. The operation (FIG. 1E) can include receiving an application event in the device-native OS (144), receiving data representative of the application event in the superstructure-dedicated OS (146),
25 applying to the superstructure, in response to the received data, a data object, thereby modifying the superstructure (148), and operating the application in the device in accordance with the modified superstructure (150). This can further include (FIG. 1F) generating a modification data object representative of the modification to be applied to the superstructure (152), translating the modification data object into a form suitable
30 for processing by the device-native OS (154), receiving in the device-native OS the

translated modification data object (156), and processing the translated modification data object in the application to update the application (158). The method can also include expressing within the superstructure a mechanism for generating the modification data object (152).

5 Still further, the step of modifying the superstructure can include (FIG. 1G) transmitting a portion of the superstructure to a processor (such as a server) remote from the device (160), modifying the transmitted portion (162), and then returning the modified portion or a new set of operations to update the superstructure (164).

10 The modification process can include using device-native code to implement an interface to modify the superstructure (148). The application of changes to the superstructure can be implemented by activating program instructions within the superstructure (148).

15 The invention can also be implemented in a network environment. For example (FIG. 1H), a copy of the superstructure can be stored on an application server operable to communicate with a remote device across a network comprising the application server, the remote device, and a communications channel therebetween (166); the superstructure can include data objects operable to instantiate applications in the remote device (166), and the network can be configured to provide communication of applications between the application server and the remote device by replicating data
20 objects in the superstructure to the remote device via the communications channel, so as to enable instantiation of new data objects and applications from the server into the remote device (168).

25 In a network implementation of the invention, when an application event is expressed to the superstructure (FIG. 1I), a superstructure object associated with the application event can be transmitted via a communications pathway from the device to a remote server (170); the server can process the object and create a new version of it, in response to the application event (172); the new version of the object can be transmitted from the server to the device to replace the existing version of the superstructure object, thus updating the superstructure (174); and the superstructure-

dedicated OS can cause the device-native OS to update the application state in response to the updated superstructure (176).

The described methods can also include providing updates to an application's state from the server to a remote device, by defining a minimal change set to the application's state and transferring it across the network from the server to the remote device, without the necessity of adapting code therefor.

Similarly, application logic can be distributed across the network by obtaining a portion of the logic from the remote device and transmitting it in a hierarchical form to the server without the necessity of adapting code therefor.

It is contemplated that the network can include a plurality of heterogeneous devices, communications channels and communications providers servicing the communications channels. (See, for example, FIGS. 4, 5, 6 and 7, which depict examples of a network using the invention, a client device, a server, and SimpleOS on a client device, respectively.) In such an instance, the superstructure can define a given application to have an appearance and behavior that can be propagated with consistency across the heterogeneous devices, communications channels and communications providers, to enable cross-device interoperability, replicability, and compatibility of applications and data with a consistency of user experience. In particular, the superstructure can be substantially free of device-specific data, modifications to the superstructure can be made in a substantially device-independent manner, and a real-time image of an application running in a first device can be expressed across the network from the first device to a second device to yield a viable instantiation of the application in the second device, regardless of device environment, wherein the organization of the superstructure and the meaning of objects within it remains substantially constant between instantiations in various device environments. Note that (as discussed in detail below) the superstructure is capable of completely expressing the running state and functionality of an application operating in a first device, and the application can be substantially identically instantiated into a second device, without loss of state or functionality, by expressing the superstructure into the second device. (It will be appreciated by those skilled in the art that the implementer

can ensure this property by maintaining the same convention of meaning for interpreting the superstructure in a given device.)

Alternatively, an application defined by the superstructure can be transmitted via a peer to peer transaction from a first device in which the application is instantiated, to
5 a second device for instantiation in the second device (FIG. 1J, 178, 179).

Among other possible variations, a method according to the invention can include validating the superstructure upon or after modification (106, FIG. 1A). If done after modifying the superstructure, the validating can include validation of data updated by processing of an event, so that the modified superstructure cannot express a
10 harmful change to the device-native OS. The method can be configured such that an application defined by the superstructure can produce external changes only by invoking operations that operate on the superstructure, thus enhancing security. (This is described in greater detail below.) The methods of the invention thus provide (among other features) an interface between an application and a system service,
15 wherein the interface is defined by interaction between the superstructure and the superstructure-dedicated OS.

Another aspect of the invention is an information processing language adapted to interface with the structures (and superstructures) described above, wherein the language (1) can be expressed entirely within the superstructure, and is capable of (2)
20 expressing a set of transformations within the superstructure, and (3) of utilizing and modifying data only within the superstructure. In this way, applications utilizing the language cannot affect the state of other applications or operate outside a bounded application container to affect an underlying device platform. This organic, structural feature further enhances security.

25 The superstructure can also contain stylesheets (910, FIG. 9) for defining selected application or presentation characteristics. These stylesheets can be configured on a per-device basis, or on a per-group-of-devices basis. Stylesheets can also be expressed within the superstructure, independent of device-specific limitations. A stylesheet can then be selected at application runtime. This helps resolve the need to
30 present a consistent look and feel for the superstructure across devices with different

characteristics, without the need to embed device-specific logic into the design of the superstructure itself.

The invention can also include (FIG. 1K) converting at least a portion of the superstructure into a device-portable form, independent of the present state of the application (180), and reconstructing the original superstructure portion, on the same or different device context, using the device portable form, without loss of state (182). This reconstructing can include using a new device-specific stylesheet. The device-portable form can be used as an intermediate or permanent storage format for recording data within the superstructure (182).

Among other variations of the methods and structures described herein, the superstructure can be organized into objects and classes (904, 906, 908, FIG. 9, described in detail below). The superstructure can also contain data structures adapted to be interpolated, and interpolation can occur whenever a device-native operating system requests data from the superstructure. The described methods can also include incorporating media assets into the superstructure, for reference by running applications, and/or incorporating by reference media assets outside the superstructure, for reference by running applications.

In another aspect, the application instantiated in the device includes an associated user interface (GUI 502, FIG. 5) whose behavior and state are defined by the superstructure. In a wireless or other handheld device, the interface can be advantageously adapted to be manipulated by a simple four-way controller, of known design, actuated by a user's thumbs (512, FIG. 5).

As described in greater detail below, the example of an SBAE operating system known as SimpleOS is adapted to run on otherwise conventional apparatus such as handheld processing devices (Personal Digital Assistants (PDAs) or others), laptop computers, desktop computers or other known microprocessor-based systems, such as those shown in FIGS. 4, 5, 6 and 7, which typically include conventional processing components such as RAM, ROM, CPU and bus elements.

For example, FIG. 4 shows a network in which the invention can operate, including a conventional server 404 of known design and function, operable to

communicate with conventional information storage elements 402 and a conventional inter-network or communications channel (such as the Internet, World Wide Web, or other) 408. Inter-network 408 is operable to communicate, in a known fashion, with client devices 410, which as shown in FIG. 4, can include (but are not limited to) a conventional cellphone, PDA, laptop computer and desktop computer (such as an IBM Personal Computer or equivalent). By way of example, the cellphone can include a conventional 4-way controller of the type having arrow keys and a central "actuate" key, as well as other thumb-actuated keys below. These physical control elements, which are shown in greater detail in FIG. 5, are used, in one practice of the invention, to provide input to a GUI that can be instantiated, maintained, modified and updated by SimpleOS, examples of which are shown and described in connection with FIGS. 11-19 below.

Similarly, FIGS. 5 and 6 show examples of a conventional client device 500 (cellphone, PDA or other device of the type represented by element 410 in FIG. 4) and a conventional server 600 (of the type represented by element 404 of FIG. 4), respectively, operable in accordance with the invention. The client device 500 of FIG. 5 is typically conventional in design and construction, including RAM 504, ROM 506, a CPU 508, a bus 510, GUI 502, and 4-way controller keys 512. In accordance with the invention, the instantiation, operation, maintenance, and updating of applications (and the GUI) can be controlled by SimpleOS 514 as described below. The interaction between the 4-way controller, SimpleOS, applications, and the GUI is also described below.

The server 600 of FIG. 6 is also typically conventional in design and construction, including such components as RAM 602, ROM 604, CPU 606, and Storage 608, collectively operable to communicate applications and/or data to and from the network.

FIG. 7 shows the relationship between SimpleOS 706 (in connection with an example of an SBAE in accordance with the invention), a device-native operating system 702, and an application 704 in a client device 700. This relationship is described in greater detail below. Similarly, in the examples of each client device shown in FIG.

4, SimpleOS is operating to receive and process applications and data 406 as described below.

II: SimpleOS Software Concepts

In one embodiment, SimpleOS is optimized for the deployment of task driven
5 (and usually, network-aware) applications with a focus on easy user interactivity,
sound, and a four-function user interface. The major goal of the SimpleOS design is to
produce a way to create a client-side application optimized for a small mobile device
that can be developed and tested on many individual device platforms without a lot of
extra work. The operating system includes many features in its design to address the
10 shortcomings of other mobile technologies and allow applications to be secure, easy to
develop and easy to distribute among a variety mobile operators.

Although the programming model set forth herein is entirely different from
previous industry designs, it does share a common feature with many platform-
independent environments in that it achieves its portability through the creation of a
15 well-defined interface. SimpleOS places an abstraction layer on top of a host operating
system such as Java, BREW or Windows that allows all operating system resources to
be consistently accessed by an application program described in XML. However,
SimpleOS does not follow the model of WAP or UIEvolution where programs are either
document-based or procedure-based. As will be demonstrated later, SimpleOS is a
20 blend between document and descriptive-based programming and procedural object-
oriented programming.

The Superstructure: A central SimpleOS feature is the concept of the
superstructure, a unifying data structure that contains the state, program code and
internal logic of each application. All of the aspects of a program that would normally be
25 controlled through private data structures, procedural code and application programmer
interfaces are incorporated into the superstructure.

Note that the superstructure can also contain data in a private form determined
by the execution environment, separate from the "public" aspects of the superstructure
viewable by the application.

Additionally, while the superstructure may be represented as a Document Object Model (DOM), its properties while running a SimpleOS-based application are significantly different. The operating system actively “watches” the superstructure as it changes in accordance with application state, and is optimized to accommodate small running changes to its hierarchical structure. Further, the operating system takes an active role in updating the superstructure and in triggering application hooks (which can also be stored within the superstructure) as a result of activity outside of the application. When a superstructure-based application is not running, or is paused, its entire state can be transferred into an XML object via a DOM or other means.

It is also significant that the SimpleOS model in part builds on an established concept in the art of XSLT transforms, where documents of one grammar are transformed to documents of a new grammar programmatically. However, SimpleOS’s use of this building block is new. SimpleOS uses a transformation language similar to XSLT to manipulate parts of a superstructure entirely within the same grammar, and provides a constant “feedback loop”, in which the output of the transformation is applied back into the superstructure causing operating system changes, and then fed back into the transformation again during the next event. To the extent these similarities are maintained in one practice of the invention, it is because they render it convenient for expressing changes to tree structures. The syntax of other languages and/or grammars could be used equivalently.

Product and Implementation: As a product, SimpleOS will appear as software that runs as a native program in a host operating system. It runs along with other provisioned applications on the device, and appears as a single application. However, the SimpleOS software is really a container for a collection of SimpleOS applications that run inside of it (i.e. an application “player” similar to a Flash player or Java Runtime Environment). The host operating system can include any C or Java based operating system including BREW, J2ME, PalmOS, WinCE, Stinger, Windows, Macintosh and Linux. To these host operating systems, SimpleOS appears as an application that can be run in the given device – for instance, a .PRC in PalmOS or an .EXE in Windows. In certain situations, SimpleOS can be configured so that applications inside of the container

appear as though they are “peers” to other non-SimpleOS applications on the device. However, when the user activates these applications, they will internally run inside a single instance of the SimpleOS native application.

Isolated, Strong Container: SimpleOS provides an isolated container that protects all running SimpleOS programs from the host environment. SimpleOS uses a strong security design that emphasizes a complete interface firewall between an application and the host environment. Within SimpleOS, a developer’s code may only access a central, well-protected document structure (the superstructure) as a means to manipulate its environment. All of the functionality that would normally be described using library or API calls is completely hidden from the application. In fact, it would be impossible for the application’s code to perform any work other than to modify the application’s own data structure. The document structures forming different SimpleOS applications are completely isolated from each other.

Consistent Interface to Multiple Platforms: The attached FIG. 8 summarizes how SimpleOS will provide a consistent interface to multiple platforms by plugging into the various API signatures of its host environment.

As shown in FIG. 8, SimpleOS implementations can be developed on a variety of mobile and desktop platforms 802, 804, 806, 808 using the C or Java language. In most implementations, the actual code of the SimpleOS product can be divided into platform-specific and platform-general code. This can enhance the potential for reuse of code within a particular programming language since different environments can swap in their own platform-specific set. For example, early versions of SimpleOS based in the C programming language contained the core functionality written in 100% pure ANSI C that did not contain any operating system calls, I/O functionality or floating-point code. A set of adapting modules allows this platform-neutral component to run on a variety of real platforms such as Win32 and Brew.

Durable Abstraction Layer: SimpleOS creates a durable abstraction layer over a host operating system that allows programs to be written in a platform-independent way, since the basic document structure that governs each application is substantially similar, not matter what platform it is running on. One practice of the invention

expresses the initial state and application code of the application using an XML grammar called SQML. SQML, like HTML, is a textual document format that allows information to be represented in a platform-independent way. However, unlike HTML, which was designed only to represent documents, SQML is designed to represent an entire SimpleOS program. Like Java, SimpleOS simplifies application development by providing a consistent platform across many host operating systems. However, unlike Java, which achieves cross-platform consistency by means of a universal machine language, SimpleOS achieves cross-platform consistency by its use of SQML.

SimpleOS introduces a paradigm shift away from the traditional programming style of programming languages and application programmer interfaces and toward a model in which the entirety of the application—the program code, application data and network communication—is completely described and operated upon using SQML. SimpleOS's brand-new model of computing is based around tree-manipulations and the complete removal of application programmer interfaces. The heart of this model is the superstructure, a data structure that resides at the core of SimpleOS processes.

It will be appreciated that other portable hierarchical and non-hierarchical structures may exist that could convey an application with the properties described herein, other than an XML structure, or a tree structure with the SQML grammar described by way of example herein. However, SQML itself, and in relation to a running SimpleOS program, enables a number of implementation-specific optimizations that allow SimpleOS to be better understood by users and more compatible with other systems or technologies that may need to interact with it. It should be noted, at any rate, that SQML could take a number of different forms, other than those illustrated below by way of example.

Superstructure Operations: The superstructure is more than just a data structure – it is the entire dynamic state of the running program, including all of its program code and data. The superstructure not only organizes the program code and data, but it also governs the execution model of the entire application. As the program runs, the superstructure evolves and changes dynamically as an exchange between the operating system and the application. The operating system translates external events, such as

user interaction, into changes to the data values within the superstructure and activates application code also located within the superstructure. When the application code reacts to a change, it responds by updating parts of its own central structure. When the operating system observes these change to the superstructure, it invokes services and library calls on behalf of the application. This intermediate step provides additional opportunities for verifying the request and optimizing it, and helps create an "interface firewall" between the application and SimpleOS. The operating system is free to choose whatever sequence or action must be taken to implement the change as it is described, freeing the program from any explicit knowledge of the particulars of the current host operating system.

The superstructure is thus the common ground between the operating system and the running application. The operating system is not required to maintain private state or data structures about the application's previous interactions since the two do not maintain any conversational interfaces such as API calls. The data that forms the specific shared assumptions between the application and the operating system exist entirely within the superstructure itself. This approach is entirely different from the traditional API or library interface model of programming that relies on application-private and operating-system private state to provide a barrier and retain the "conversation." Nothing in the present invention, however, prevents the ability to maintain private state if desired or useful.

Tree Structure Manipulations: Because the superstructure is a tree, application code preferably can be restricted only to operations that involve manipulating trees. All outside data is delivered in the form of a tree, and external changes to an application occur exclusively as a result of tree operations. Consequently, application code is based entirely around tree transformations, and the work done by the application is *descriptive*. There is no explicit application programmer interface (API) needed for SimpleOS programs. However, an API can still be employed if desired, and the nature of the present invention allows the possibility of using an API. For interactions that are based more on "verbs" as opposed to descriptions, such as low-level network or

physical events (such as ejecting a disk), the descriptive approach may be less compelling.

Traditional Models vs. SimpleOS

In the traditional programming model, the programmer writes code using a

5 programming language that is compiled into the native instruction set or machine code of the target platform. In the case of byte-code interpreted languages like Java or Perl, the machine code that results from compilation is not the machine code of the native hardware, and further translation or interpreting is necessary for execution of the code. The programming language itself typically only supports program decision-making,
10 mathematical operations, control-flow and other computing fundamentals. In order to perform anything beyond pure computation, the programmer uses a platform-specific library of external Application Programmer Interface (API) calls to perform functions not provided by the programming language. These functions typically include such operating system functions as file input/output, memory management and process
15 control, as well as functions provided by a program library, for example a graphics library. Therefore, within the traditional programming model there is an abstraction barrier between the application and the operating system based on the operating system's API.

Traditionally, the operating system's API acts as an intermediary between the
20 state and assumptions of the operating system itself and the internal state the application. The operating system's own data is hidden from the application, often in a protected area of memory. The API is the only means for the application to change the state of the operating system. The API may be thought of as a set of services that the operating system provides to its applications.

25 In existing programming environments, the API is needed to enforce the separation of state between the assumptions and data space of the application and its host operating system. Typically, the application manages its own data structures by dividing its binary address space into areas implicitly separated by the weaving of the program code, stack frames, heap allocations, etc. Typically the programmer crafts

useful abstractions inside this private memory area to help in the operation of the program and keep track of its conversational state with the operating system. The operating system is completely ignorant of the private data structures of the application. Similarly, the operating system manages its tasks by creating private data structures to represent outside interaction and providing API calls for the application to modify these OS private structures as part of its conversational state.

Note that in referring to "conversational state" herein, the applicants refer to the back-and-forth API calls between a traditional application and a traditional operating system. By way of example, consider opening a file and writing to it. In such a case, the application asks for a file handle from the operating system. The OS then allocates some private structures to account for the transaction on its end, and then passes something back to the application, called a file-pointer or stream handle, or something that the application then private stores. When the application writes to the file, it must pull up that value or number, and use it to reference the file with the OS. The present invention, in contrast, is much more like memory mapped IO, in which the file itself is reflected into the superstructure. The application asks for the file to be opened by setting up a node in its superstructure with a filename. The operating system sees this new node, and maps the contents of the file into the superstructure. The application then reads from it, just as if the entire file were part of the superstructure. The advantage of this is that using well-known techniques such as paging and buffers, the operating system can very efficiently maintain this virtual structure (i.e., that the entire file is part of the superstructure).

In addition, as schematically shown in FIG. 3, the operating system can use the internal representation of the superstructure (i.e., the portion that the application does not "see") as a hook for placing private data related to the request (304, FIG. 3). The superstructure must internally be represented as a set of objects with references, a set of structures with pointers or the like. Either way, there is a private representation of the superstructure that the SimpleOS implementation uses to store in the host's architecture. The noted structure can be used to store data beyond just the contents of the superstructure. It could also store private information related to the particular state

or data type of the node. For example, the file-opening node discussed above could be represented as a data structure that has a "spare field", readable only by the OS, that contains the file pointer provided to SimpleOS from the host operating system. This private data is not serialized when the application is paused, halted or migrated, but it provides a convenient place to store this information that is precisely organized in an efficient manner for the SimpleOS implementation while the implementation is running (304, FIG. 3). The result is an ability to store non-conversational data in a completely recoverable way, so long as the private data can be re-established upon deserialization based on public data that has been maintained in the superstructure. This is primarily for efficiency.

SimpleOS does not require an API because it does not allow the application to explicitly invoke its services. Instead, the SimpleOS operating system acts as a constant observer to the superstructure. When the application changes its superstructure, the operating system automatically invokes the services that are needed to perform the change in the outside world. When outside events occur, the operating system passes them down into the superstructure by manipulating data values inside and providing notification to the application. The superstructure completely replaces the explicit invocation of outside resources in an application.

Another way to view the difference between SimpleOS and traditional models is the structure of the code. In the traditional model, the program is *prescriptive*. Like a recipe, it provides a list of instructions to the operating system that prescribe a particular activity, such as "draw a dialog box here, then add a button, then render some text." In the SimpleOS model, code is *descriptive*. The SimpleOS application describes the final state of what it would like to happen ("a dialog box on the screen, with a button and some text"). The operating system takes this description and determines the optimal way to implement the change.

Using the superstructure concept described herein, the invention in effect provides a way to adapt the paradigm of memory mapped IO to accomplish substantially every task within an operating system, not merely file writing.

Contents of the Superstructure

The superstructure includes the entire nature of an application program. Application program code, private and public data, graphical user interface configuration and virtually every single aspect of a program is kept together in a single, highly organized data structure containing the shape of a tree. An application's permanent and running state is entirely defined by the superstructure, including:

- Every detail about the user interface, including colors, text, positioning, user elements, screens, backgrounds, sounds and controls
- All internal application state, including local data, state variables
- All client-side application code, in the form of message registrations and transformations

The composition of the superstructure is described in detail in "Composition of the Superstructure" below. For now, it is important to know that the superstructure is internally defined as a single, very extensive tree that contains objects that describe the state of the application and its interaction with the operating system. Each node inside the superstructure can contain attributes and children nodes.

Programming Paradigm

The XML source of a SimpleOS application is not a program structure that is interpreted or compiled. Instead, it can be the seed image of the initial state of the superstructure. When a program is started, this seed image populates a new superstructure for the application.

The state of the superstructure at any point in the execution of a SimpleOS process may be represented in SQLML. An application's state may therefore be captured in SQLML form and transmitted or saved to a file. Or, the state of the superstructure may be represented in a more compact, platform-independent binary format. This allows a SimpleOS program to be saved and resumed, or to migrate from one system to another, without loss of continuity. Because of the equivalence of the superstructure to SQLML, a running application process can be easily condensed into a serial form that can be communicated across a network or other streaming connection.

Again, SQML, as described herein, is but one grammar that could be used to seed a superstructure. In addition, while the applicants describe herein particular examples of organizing the superstructure in SimpleOS, many different organizations could be used to provide a useful implementation based on the broad concepts of the invention, and are within the scope of the present invention.

Dynamic Sections

The actual organization of the superstructure maintains the abstraction between the operating system and the application. A defined contract between SimpleOS and the application maintains the meaning of each section of the superstructure based on well-defined names, types and locations. In the illustrated embodiments, this contract is expressed as part of the SQML grammar and the specification of SimpleOS itself. This could of course be different in different implementations of the invention. The operating system's work is directly affected by the content of particular sections of the superstructure, and any changes to those sections of the superstructure effect an immediate change within the outside application. No API calls are required by the application to read or write data from the operating system. Instead, all operating system functions are accessed by changing these "dynamic sections" of the tree that share common meaning between the operating system and the application. Again, the overarching principle is that the contract of the dynamic sections relates to a set of shared semantic assumptions between the application and the operating system about the meaning of each section of the tree. The specifics of implementation, however, can vary in accordance with the implementer's design.

In essence, when an application wants to manipulate its external interface, it describes changes to particular parts of its superstructure. The operating system observes these changes to the superstructure, and invokes services on behalf of the application to perform the work described there. The operating system is free to choose whatever sequence or action must be taken to implement the change as it is described, freeing the program from any explicit knowledge of the particulars of the current host operating system.

The term “dynamic sections” is used herein to describe the elements within the superstructure that have shared meaning between the operating system and the application. These elements are constantly observed by the operating system for change. The “dynamic” sections of the tree are defined by a grammar that indicates the node types, locations and attributes. The parts of the application that have “dynamic” state include the nodes that represent the graphical user interface, the event registration, input and output, etc.

To better appreciate the implications of this organization, consider a simple example. For instance, assume a hypothetical mapping wherein a node within the superstructure named “root.gui.screen.mainScreen.myButton” represents a button that an application displays on the screen at startup time. The startup SQL of this application describes a node with this pathname and attributes that describe the properties of this button, including its location, size and color. For instance, if the application initially wants the button to have a blue color, it will define an attribute on this node with the name “foregroundColor” that contains a value of “blue”. The very act of including this node within the superstructure will immediately cause a button to appear on the screen when the application starts – no additional work is needed. Because the node exists in a “dynamic section” of the superstructure, it has meaning to both the operating system and the application, and both of them maintain its meaning through the lifetime of the node.

Continuing the example, assume that the application would like to change the color of the button from blue to red. The application merely has to update the same GUI node and establish a different value for the “foregroundColor” attribute. Immediately after this change has been made, the button’s color will change on the screen. The reason is that in the SimpleOS execution model, the operating system efficiently tracks the changes made to the superstructure after every code fragment has run. Upon discovering that the code has changed an attribute on a button, the operating system will verify the change and call the appropriate services on the device to actually change the visible color.

Describing Programs with SQScript

In the SimpleOS technology, to manipulate the superstructure, the application program describes modifications to the tree using a rich expression language called **SQScript**.

SQScript is a powerful language that is similar in syntax to XSLT, but much more
5 dynamic. Within SQScript it is possible to change data values within the superstructure, including those in dynamic sections, using tree modifications and updates. As noted herein, SQScript is but one of many languages that could be used to implement the present invention. Other XML or tree-structure-based languages capable of
10 implementing private variables, program control, event handling and tree manipulation could also be used.

Other Programming Approaches

As described later, the application may also write its code using native languages on the device or over the network. See "Programming SimpleOS" below.

Application Lifecycle

15 As noted above, SimpleOS applications may be provisioned from an XML grammar that describes the initial state of the superstructure. In a SimpleOS implementation of the invention, each application stored within the SimpleOS container has exactly one initial state definition. Each provisioned application also has a unique ID that is used to manage consistent network communication. Applications typically reside in a permanent
20 area of memory until the user activates them. This does not necessarily mean that all of the resources for the application have been loaded. If desirable, the operating system can purge infrequently used application resources from the device and re-load them over the network if they are needed.

At any point in time, a running application has one or more executing processes.

25 Each process has its own copy of the application superstructure, and maintains a message queue of pending operations. If the developer wishes, the application can be started as many times as the user wishes, so that a given application can have more

then one running process. SimpleOS provides the user with a mechanism to switch between running processes.

In the reference implementation of SimpleOS (which is but one possible implementation of the invention), each process runs in its own thread, allowing true multitasking between running processes. However, it is not necessary for the host platform to support multi-threading. Since many of the target platforms such as Brew lack any support for threads, the SimpleOS implementation on these platforms must provide an internal form of multitasking sufficient for consistent execution behavior across all platforms.

10 ***Developing Applications***

In a particularly useful implementation of the invention, applications are provisioned to the operating system by conveying or transmitting a "seed" image of the application in some portable format. Thus, for example, developing a SimpleOS application involves writing the "seed" image of the application superstructure. Outside the device, this seed image is typically expressed in SQML, which is an XML grammar that represents any type of data inside the superstructure. SQML (or whatever appropriate grammar is used to convey the superstructure) is not itself a programming language like C or Java.

Instead, it should be thought of as a combination of the initial state of an application and the rules that show how the initial state reacts to the actions of the user. Under this paradigm, development is akin to filling in a large tree that organizes different aspects of the application, clustered into objects and classes. The developer specifies screens and controls by grouping them into classes that are stored in the superstructure. In the SimpleOS examples discussed below, the actual code for each handler is written using a template and expression language called SQScript.

25 ***Running Applications***

SimpleOS (or a similar superstructure application environment constructed in accordance with the invention) comprises a strong container implemented on top of an existing host operating system. In the most basic case, SimpleOS will appear as part of

a menu of applications available on a device, for instance as an icon on a Palm's launcher application or an item in the Window's Start menu. Once the SimpleOS container is started, it will automatically invoke a special "launcher" application, written in SQLML and native platform code. The launcher application will query the internal storage on the device and present a menu of available applications. The user can browse the list of applications that have been already provisioned and choose one to launch. At this point, the seed image of the new application is placed in its own process, populated into a new superstructure and launched. The application will take over the screen, initialize its functionality, and begin to receive messages from SimpleOS.

The launcher application will also support a network mode, where the launcher can download lists of new applications and trigger the provisioning process. Applications can even be browsed "on-the-fly" by entering a URL. In this case, they are never formally provisioned to the device but reside in temporary memory. This functionality allows SimpleOS (or a similar superstructure-based technology constructed in accordance with the invention) to support a "WAP-style" of connected application browsing.

The launcher is a hybrid application, written partly in SQLML and partly in native code. It requires native support in order to perform functions normally disallowed by the operating system such as process and application control.

Provisioning Applications

SimpleOS will store a number of SQLML applications in persistent storage to reduce bandwidth and latency for launching applications. The manufacturer or the user can provision the majority of applications, in an effort to contain bandwidth and improve startup time. Mobile Operators or Simple Quick GO will be able to determine a menu of applications available to a particular user that can be downloaded over the network itself.

The technological approach of SimpleOS (or other implementations in accordance with the invention) supports several desirable provisioning models for wireless handsets and other mobile devices. The key to the ease of provisioning is the

ability for the application "source", which is really just the seed image of the superstructure, to be expressed as SQML (or an equivalent transport mechanism). In its streamed form, such a document can easily travel over any stream-based connection protocol including raw TCP/IP, HTTP, WSP, and serial connections. Several of these protocols, including HTTP and WSP, are supported over wireless networks such as cellular, 802.11b (WiFi), and Bluetooth. Planned implementations of SimpleOS will contain support for a variety of these protocols. For simplicity, a present implementation of SimpleOS relies on HTTP and HTTPS, which may be layered upon other wireless or wireline protocols.

Typically an application will consist of SQML and some number of media assets such as images, sounds, graphics, animations, videos and MIDI tunes. Standards such as Internet MIME can be used to lace together these different media files along with the SQML document into a single application image that can be shipped to the device using network protocols.

Server Provisioning

In one embodiment, the present invention can provide a J2EE-based provisioning server to download applications to the device upon request. The provisioning server will maintain a registry of all of the devices that can receive applications from it, including their security keys and their device characteristics. The provisioning server allows for applications and data to be transmitted securely to each terminal and for the application data to be pre-formatted to fit the capabilities of the particular device. By knowing the characteristics of the device, images can be resized, sounds files can be converted and media stream tags can be changed based on the device capability.

When the launcher application makes a request to the server for a new application, it will include its provisioning key with the request. This provisioning key is a unique identifier for the device that was set either at the factory or during the user's installation of SimpleOS. The provisioning server will perform an access check to ensure that the provisioning key is valid, and then look up the device characteristics. Now armed with complete knowledge of the device capabilities, the provisioning server can

perform a variety of up-front tasks to improve the application on the device. This may include:

- Resizing images or adding/removing colors to suit display capabilities
- Changing the compression format of images, sounds and media to match de-compressors found on the device
- Altering the bit-rate of sound and media formats to avoid overwhelming the device
- Pre-applying style-sheets or conditional macros within program code

The resulting SQML and associated media assets are laced together into a single application file and returned to the device. The conversion activity may occur on the fly when the request is issued, but for speed the model also supports the provisioning activities as a one-time operation done in advance for each known device type.

On the device, the SQML and other assets will be filed inside the internal file storage available to the SimpleOS container. The application will be registered internally so that it appears within the launcher with the appropriate icon and label.

As noted above, methods other than SQML can be used for conveying and launching a superstructure-based application in accordance with the invention.

Peer Provisioning

One embodiment of SimpleOS also supports a provisioning model where applications are sent from peers over protocols such as 802.11b, Bluetooth, and Infrared. In this case, the launcher application on one device can send a message to the launcher application on another device with the content of the application. This method would also work for provisioning kiosks, where a user simply places his device near or on the provisioning kiosk and software within the kiosk initiates a transfer of an application to the device.

III: Composition of the Superstructure

In accordance with one embodiment of the invention, the superstructure represents the entire nature of a SimpleOS application program. Application program code, private and

public data, graphical user interface configuration and virtually every single aspect of a program is kept together in this single, highly organized data structure. When an application is launched, a new superstructure is created based on the SQML definition (seed) of the program.

- 5 Conceptually, the superstructure has a tree shape that loosely follows the conventions of an XML Document Object Model (DOM), where each node is like an XML element. Each node/element may have any number of attributes characterized by a name and string value. Nodes can be nested according to the grammar of the SQML DTD, a specialized grammar that governs the permissible shapes of the superstructure.
- 10 It should be noted that the SQML DTD is not the grammar of the SQScript language – it is actually the grammar of the continually updated data structure of the running application.

- However, the superstructure of the present invention differs from an XML DOM. Unlike an XML DOM, the superstructure inherently allows groupings of entities into
- 15 objects and classes. And also unlike an XML DOM, the superstructure of the invention regards these grouped entity objects to have methods, inheritance and classes that affect their data retrieval and execution properties.

- For purposes of simplicity, in the examples discussed herein, the applicants have opted to use a number of XML standards as part of the SimpleOS implementation.
- 20 However, those skilled in the art will understand that the properties of a superstructure-based application programming environment like that of the invention transcend specific XML grammars or approaches. In particular, a superstructure-based environment utilizes a hierarchical structure with a specific and controlled grammar to provide a consistent structure and a way to express the contract between the operating system
- 25 and the application. XML-based technologies share some of these properties and can be used to implement the present invention. However, the superstructure-based application programming model of the invention could also be implemented using other approaches.

- It should also be noted that in the interest of simplicity and syntactic ease of
- 30 programming, the SimpleOS implementation of the invention described herein does not

use XML as the grammar in all instances of its programming. Again, one seeking to implement the invention could do so in various equivalent ways.

The following sections describe a SimpleOS implementation of the invention, utilizing a series of conventions, names, structures and parameters specific to particular examples of a superstructure-based application environment (SBAE) in accordance with the invention. Those skilled in the art will appreciate, however, that these conventions are a matter of design choice, and were selected to make the illustrated implementation simpler and more robust. Again, many variations are possible and within the scope of the present invention.

Top-Level Branches

Conceptually, data within the superstructure is divided into three “branches” as shown in FIG. 9: Stylesheets 910, Classes 908 and Objects 904, 906. Most of the application’s data is stored in objects, which can be created from classes. In the illustrated embodiments, stylesheets govern the visual presentation of the SimpleOS user interface. The following sections refer to the 3 major objects types shown in FIG. 9.

Objects and Classes

Directly under the root of the superstructure are a number of *object instances*. Each object instance has a unique name called an *instance name*. No two objects in the process’s superstructure may possess the same instance name. In the examples shown, each object contains structured data that may include the following types of data defined by the SQML DTD:

- **Methods:** Code definitions that describe a dynamic modification to the object
- **Data Members:** The equivalent of variables, which can be of any of the supported data types.
- **GUI Definitions:** Definitions for cards, frame buttons and other graphical elements.
- Additional types, defined in the future

Within the programming language, new objects are created by process similar to a data structure clone. This operation either clones a pre-existing object or creates an object based on a special template called a *class*. A class definition can be identical to an object's definition, except that it can only serve as a template. Any screens, or other data types stored within the class are unavailable unless they are used to create object instances.

Each object is basically a copy of the class; the class is purely a template. Every object has as parts of its structure the original class ID, so it remembers the class type. When an object is cloned from another object, it doesn't remember what the object ID of the original object however it retains the original class ID no matter how many times the object is cloned.

In one practice of the invention, within SQML, both objects and classes are defined using the sq-class tag. A class is defined using syntax such as:

```
<sql:sq-class classid="class_menu" extends="class_menubase">  
  <sql:variable name="index">1</sql:variable>  
</sql:sq-class>
```

An object instance is defined using an identical syntax except that an objectId attribute is included in the sq-class element. The following represents an object called main:

```
<sql:sq-class extends="class_main" objectid="main">  
  <sql:variable name="index">123</sql:variable>  
</sql:sq-class>
```

The example above actually defines both a class called class_main and an object called main. Both class_main and main could be used as the source of a new object instantiation. In the case of instantiating an object from class_main, you would be creating a new instance of an object of the class class_main. Instantiating the object using main as a template would cause main to be cloned.

Stylesheets

Besides classes and objects, there is a third type branch of data within the superstructure called a *stylesheet*. A stylesheet defines a set of visual formatting characteristics that are used by the user interface presentation code. Stylesheets are discussed in great detail below. Table 1 shows a example of a stylesheet.

The implementation of stylesheets, per se, is well-known in the art. However, their use in a superstructure-based application environment is an advantageous and novel aspect of the present invention, representing a useful and efficient way of making a platform-independent architecture have a meaningful and expressive appearance within each specific implementing architecture.

```
<sqlml:stylesheet>
<sqlml:class id="base">
  <sqlml:palette>
    <sqlml:colordef name="white" rgb="#ffff00"/>
    <sqlml:colordef name="black" rgb="#ffff00"/>
    <sqlml:colordef name="red" rgb="#cc0000"/>
    <sqlml:colordef name="blue" rgb="#4444ff"/>
    <sqlml:colordef name="yellow" rgb="#ffff00"/>
  </sqlml:palette>
</sqlml:class>
<sqlml:class super="base" default="true">
  <sqlml:bgcolor color="white"/>
  <sqlml:font typeface="Helvetica" typesize="12" typecolor="black"/>
</sqlml:class>
<sqlml:class super="base" id="border">
  <sqlml:bgcolor color="blue"/>
  <sqlml:font typeface="Helvetica" typesize="16" typecolor="white"
    typestyle="bold" />
</sqlml:class>
</sqlml:stylesheet>
```

Table 1 - A Stylesheet Example

Object Overview

Objects group structured data, variables and methods together in a single named entity. As the highest-level branches within the superstructure, they maintain the basic organization of a running program. There are three types entities generally found within objects:

1. Object Data Trees, containing arbitrary XML structured data

2. Object Data Members, data types available within the SQScript programming language
3. Object Methods, SQScript functions that can be invoked on the object.

5 FIG. 10 summarizes the proceeding sections on each of these three types, showing object data trees 1014, object data members 1016 and methods 1018.

Data Trees

One of the primary functions of an object is to provide a way of grouping SQML elements together into a “miniature DOM” so that it can be referenced and accessed
 10 during the operation of a SimpleOS program. In the illustrated embodiments this data is typically used to describe graphical user interface elements, although the data may be any valid type of data expressible with SQML. The object’s XML data payload is called the object’s *data tree*.

Table 2 demonstrates an object that contains an XML data type defined in SQML
 15 called a sq-card. The sq-card definition is embedded inside the XML that described the object. Inside the superstructure, this sq-card definition will be attached underneath the object “main”.

```

20 <sql:sq-class classId="class_main" objectId="main">
    <!-- The main card -->
    <sql:sq-card id="card_main">
      <sql:title label="MONOIDS" />
      <sql:top label="ABOUT US" class="border" />
25 <sql:left label="INSTRUCTIONS" class="border" />
      <sql:right label="MENU" class="border" />
      <sql:bottom label="MENU" class="border" />
      <sql:leftButton label="MENU" />
      <sql:rightButton label="EXIT" />
30 <sql:frame>
      <sql:image src="splash.gif"/>
      </sql:frame>
    </sql:sq-card>
35 </sql>
  
```

Table 2 - A Simple Object with a Data Tree

Element Identifiers

Each element inside the data tree of an object may contain an id attribute. For instance, the sq-card definition in Table 2 has an id attribute set to card_main. The id tags of the elements in a single object must be completely unique within that object. For instance, while it would be possible for an object called differentObject to also contain an element named card_main, the original main object may not have additional elements with the id card_name.

Because of the uniqueness constraint over the elements within an object's data tree, and the uniqueness constraint over the name of an object, a special namespace exists if these names are combined. This naming convention is called an *element identifier*. In one practice of the invention, any element inside the data tree of any object may be uniquely addressed by the element identifier syntax

objectID#elementId where objectID refers to the ID of an object and elementID refers to the ID of an element within the object. As a special rule, the objectID portion of the identifier may be omitted within the same object.

Table 3 demonstrates several different formats for referring to elements within a hypothetical superstructure.

main#card_main	The element "card_main" within an object called main.
myObject#theButton	The element "theButton" within the object myObject.
#nextScreen	The element "nextScreen" within the current object.
myObject#	Illegal – no element ID is present.

Table 3 - Examples of Identifiers

Dynamic Sections Revisited

Data tree information may be part of a dynamic section, particularly if it describes GUI data. If that is the case, any changes to the model represented by the structured data within an object will automatically be reflected outside the application. For example, referring to Table 3, assume that main#card_main is the currently active screen. If any part of this definition within the object is updated, such as the name of a frame side,

the physical display on the device will also be updated. Similarly, if the user updates the contents of a data field on the screen, that change will immediately cause the relevant attribute within the object to change.

Data Members

Data members form a special class of elements that may be found within the data tree of an object or class. A data member is part of the element data within an object, but may be accessed in the SimpleOS programming language SQScript as a local variable. This feature is covered in detail in "Programming SimpleOS Applications" below.

Table 4 demonstrates a set of data members defined on a class called `class_menubase`.

Data members may be freely mixed with other valid data types, although only the list and scalar syntaxes may be referenced as variables. In one practice of the invention, expressions found within variable definitions are evaluated when the class is instantiated.

```
<sql:sq-class id="class_menubase">
  <sql:list name="list_pizzaNames">
    <sql:list-element>thin crust pizza</sql:list-element>
    <sql:list-element>thick crust pizza</sql:list-element>
    <sql:list-element>deep dish pizza</sql:list-element>
  </sql:list>
  <sql:variable name="foo">One Two Three</sql:variable>
  <sql:variable name="bar">1232</sql:scalar>
</sql:sq-class>
```

Table 4 - Examples of Object Data Members

Methods

An object may also contain methods, which are custom-written programming functions written in SQScript. (See "Programming SimpleOS Applications" below.) A method definition uses the function element tag within an object's data tree. Table 5 shows a simple example of the use of this syntax. It will be understood that while

superstructure applications could be implemented without the use of methods, members and the like, they represent highly useful features of a SimpleOS implementation of the invention.

```
5      <function id="next">
        <sql:eval expr="index=$index+1" />
        <sql:if test="$index > $list_pizzaNames:size">
          <sql:eval expr="$index=1" />
        </sql:if>
10    </function>
```

Table 5 - Defining Methods for an Object

Graphical User Interface Data Trees

15 One frequently used data type stored within the data tree of an object is the definition of GUI elements that are rendered by the **Simple Quick Go Graphical User Interface** (SQUI). In typical implementations, definitions for GUI elements such as SQCards will be among the most common type of data stored within an object's data tree. A complex data grammar governs GUI definitions, with structures that parallel the
20 visual presentation of the display.

Visual Presentation

To help explain the user interface sections of the SimpleOS superstructure, the following section describes examples of a SimpleOS user interface.

25 In one embodiment, the user interface is divided into a number of nested elements such as screens, controls, and images. Each of these elements has a *style*, which is a reference to a stylesheet class that defines the presentation attributes of the element: color, font, and layout. No specifics about the presentation of an element are contained within the GUI description area of the superstructure. (See "Stylesheets" below.)

30 Because of the limited screen size available on most implementations, the graphical user interface is divided up into a series of screens called "SQCards". Only one SQCard may be displayed at a time. The operating system allows the developer to create a number of different screens at once, and choose which one should be

displayed. Typically the developer will organize specific areas of application functionality by screen. In some cases, one screen flows into another naturally by selecting a control or pressing a frame button. In other cases, a screen can indicate an unexpected situation such as an alert.

5 Each SQCard contains a number of attributes and elements. A SQCard must have a unique element id within its class. A SQCard also contains a title, either as an image or as a text screen. There are other style options available, such as a background image, spacing parameters and scrollbars.

10 Each SQCard's display includes four labels at the top, left, right and bottom of the screen, forming a central square frame. These are the **frame buttons**. They correspond visually to the buttons on the device's four-way controller. Within the region bounded by these labels appear the main content of the SQCard. The square frame formed by the four frame buttons is called the **action frame**.

15 Consider FIG. 11, showing the visual presentation of a GUI 1102. The look is dominated by the action frame 1120. Immediately above the action frame is a **title area** 1104, which can either be a graphics image or a text label 1108. This label tells the user the current name of the screen. Above the title label is a **status area** 1106 that contains information about the device, such as the date and time, and the online/offline status of the application. Below the action frame is a **soft key area** 1122
20 where labels that correspond to physical buttons on the device are shown.

 The area inside the action frame is considered the **control region 1124**. Additional controls and content are presented inside this area. Controls can include checkboxes, text labels, radio buttons, images, icons, and other typical, known user interface components. If the items inside the control region exceed the size available,
25 **scroll bars** 1110 will become visible that show the relative area of the virtual control area that is visible on the screen.

GUI Data Examples

By way of example, the following SQML (Table 6) defines a very basic SQCard called "card00". This code can be inserted as part of an object's data tree, where the GUI engine can later find it if the screen becomes active.

```
5  <sq-card id="card00">
    <top label="MOVE UP" id="top" class="blue-on-yellow"/>
    <left label="MOVE LEFT" id="left"/>
    <right label="MOVE RIGHT" id="right" class="blue-on-yellow"/>
10  <bottom label="MOVE DOWN" id="bottom"/>
    </sq-card>
```

Table 6 - Sample Definition of an SQCard

15 The main control region inside the action frame is defined by a frame tag. This element may include checkboxes, radio buttons, text areas, images and other common graphical widgets.

 The following example (Table 7) shows the definition of a frame inside an
20 SQCard. Each editable area of the screen contains a value attribute that is updated automatically when the user operates the interface and changes information inside controls. These value attributes live within the superstructure itself, and can be accessed using the normal means for looking up data values.

```
25  <frame class="plain" id="frame00">
    <image id="comp0" src="icon1.gif"/>
    <checkbox id="comp1" label="this is a checkbox"/>
    <radiobutton id="comp2" label="this is a
30      radiobutton"/>
    <textarea class="blue-on-yellow" id="comp3">
        Some text
    </textarea>
35  </frame>
```

Table 7 - Sample Frame Definition

The “Active” Card

In one practice of the invention, at any given time, there is always one active sq-card definition, found within an instantiated object within the superstructure. This means that the most basic application possible must contain at least one object instance that has at least one sq-card definition within its data tree. A variable within the dynamic section of the superstructure, called ui#card tracks the instance name and element id of the current sq-card. This value is kept as an element identifier string, and stored on an attribute found within the *implicit* UI object. For more information about the implicit UI object, see "Implicit Objects" below.

As a convenience to the programmer, the initial value of the active screen variable can be set using the “activate” attribute on the sqml tag. See the example shown in Table 8.

```
<sqml appid="com.sqgo.simpleos.application.demo"
      activate="main#card00">
```

Table 8 – Example: Defining the Initial Card

Example Static Application

The following sample application demonstrates what a simple application might look like when represented as SQLML. (Again, many variations are possible, and the invention does not mandate the structure shown by way of example.) This application presents a simple screen to the user that looks like that shown in FIG. 12.

The source of the application depicted in FIG. 12 is shown in Table 9. Note that stylesheets are used in this example. Stylesheets are discussed in detail below.

The basic structure of the application starts with a sqml tag, which represents the root of the superstructure. Below it is defined one object called main, which contains a sq-card definition inside its data tree. The activate tag on the sqml tag causes this card to be initially displayed to the user. See Table 9.

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE sqml SYSTEM "sqml.dtd">
```

```

<!--
  A simple application without any dynamic handlers
  Demonstrating controls and cards.
-->
5  <sql appid="com.sqgo.simpleos.application.demo"
    activate="main:card00">

    <sq-style-sheet>
10      <class id="base">
          <palette>
              <colordef name="blue" rgb="#4444ff"/>
              <colordef name="yellow" rgb="#ffff00"/>
              <colordef name="white" rgb="#ffff00"/>
15      <colordef name="pink" rgb="#ff8888"/>
          </palette>
        </class>
        <class super="base" default="true">
            <bgcolor color="blue"/>
20      <font typeface="Courier" typesize="14" typecolor="white"
            typestyle="bold" />
        </class>
        <class super="base" id="blue-on-yellow">
            <bgcolor color="yeller"/>
25      <font typesize="10" typecolor="blue" />
        </class>
        <class super="base" id="plain">
            <bgcolor color="pink"/>
        </class>
30    </sq-style-sheet>

    <sq-class classid="class_main" objectid="main">

        <sq-card id="card00">
35      <top label="MOVE UP" id="top" class="blue-on-yellow"/>
          <left label="MOVE LEFT" id="left"/>
          <right label="MOVE RIGHT" id="right" class="blue-on-
          yellow"/>
          <bottom label="MOVE DOWN" id="bottom"/>
40      <frame class="plain" id="frame00">
          <image id="comp0" src="icon1.gif"/>
          <checkbox id="comp1" label="this is a checkbox"/>
          <radiobutton id="comp2" label="this is a
45      radiobutton"/>

          <textarea class="blue-on-yellow" id="comp3">

50      Here is a text area that should spread onto
          multiple lines. We the people, in order to
          form a more perfect union, establish justice,
          ensure domestic tranquility, provide for the
          common defense, promote the general welfare and
          secure the blessings of liberty for ourselves
          and our posterity do ordain and establish this
55      Constitution of the United States of America.

```

```

5      </textarea>
      </frame>
      </sq-card>
    </sq-class>
  </sql>

```

Table 9 - Static SimpleOS Application Example in SQLML

Implicit Objects

One of the key design principals of SimpleOS is that the application superstructure represents the entire state of the running program at all times. In order to ensure this is always the case, SimpleOS (or similar system in accordance with the invention) may define a number of *implicit objects*. (Note that while it is possible to configure a superstructure-based application environment without implicit objects, they are of great utility in configuring a useful result.) These objects represent the state of the operating system and are essentially inherent "dynamic sections." For example, the ui object that described above represents the state of the visual display, which belongs to the operating system rather than to the application. These implicit objects may or may not be accessible to the application via SQScript.

An implicit object behaves as a regular object within the superstructure. It contains data attributes, methods and a structure data tree. However, these objects are always found in the superstructure, whether or not the seed SQLML image for the application defines them. They are created and populated when the SimpleOS loader first creates the superstructure for a new process.

If the SQLML source code of an application does define an implicit object, the values found within the definition will be used to populate the contents of the implicit object. It is impossible to remove an implicit object from memory at any time.

Implicit objects are almost always part of the dynamic section of the superstructure, since they typically reflect meaning that varies from process to process and changes over time. Implicit object can be used by SimpleOS to represent the state of the user interface, other applications in memory, the queue of currently playing sounds and external interfaces such as Bluetooth and 802.11b.

IV: Programming SimpleOS Applications

In accordance with the principles of the present invention, the SimpleOS embodiments described herein (and similar configurations in accordance with the invention) allow for an implementation in which basic types of applications can be written entirely through the use of a drag-and-drop editor, thus greatly simplifying the creation of applications. More complex applications may be written using both server-based and client-based code. The following section describes examples of various ways to program SimpleOS applications (SimpleOS being but one example of the invention), including network and client-local programming.

Overview

Changes inside of a SimpleOS application are simply updates to a tree – usually, the data tree within a SimpleOS object. This tree update may include changes to a basic attribute on an existing node or may involve widespread changes to the entire organization of the object in question. Programs to modify the current tree may be written using a variety of methods. They may include (but are not limited to):

- client-native,
- built into SimpleOS,
- written using SQScript, a platform independent programming language, or
- written using any common server-side programming language.

It is a significant feature of the invention that while the various programming models are very different, the superstructure modifications performed by the programs are the same. This focus on tree manipulation calls for a programming model that is optimized for tree transformations.

The superstructure-based model of programming shifts almost all of the usual work a program does to maintain its user interface and other external interfaces into modifications of a tree. SimpleOS applications do not need to concern themselves with maintaining communication with an Application Programmer Interface or with the nuances of user interaction. All of the general environmental and user-interaction

functionality is handled for the application by SimpleOS. When application code is run, all of the state of the application has already been updated inside the superstructure. The code merely has to change data within the superstructure, and the complex work of making those changes apparent within the host operating system occurs
5 automatically.

There is very little difference between programming on a network server or directly on the device. Data within the superstructure can be efficiently communicated over a network, and a set of tree-updates can easily be communicated back to the device. Consequently, there is almost no friction for a network program that wishes to
10 perform equivalent operations to client-local code.

Core Programming Modes

It is expected that many implementations of the invention will naturally gravitate toward having much of the application logic expressed as some local programming embedded therein. However, other mechanisms, including those described above, can
15 also be supported. SimpleOS, among examples of the invention set forth herein, supports at least four basic mechanisms that illustrate how client-native, superstructure-local, network, or built-in approaches are enabled.

In a SimpleOS practice of the invention, a highly useful model for coding is a special programming language called SQScript. SQScript is optimized for tree
20 transformations and can run on any implementation of SimpleOS. There are also several alternative methods besides SQScript. The application can use native-language code to manipulate the superstructure through the use of a special API. In this case, any code that can be dynamically or statically linked into the SimpleOS container can be used to manipulate the superstructure. Of course, using native code eliminates many of
25 the advantages of portability and security gained by SQScript. Another alternative is to use a set of "default handlers" that provide common pieces of application functionality out of the box. Finally, code can reside on the network, allowing a server to provide all of the application logic.

Each of the four approaches to SimpleOS programming revolve around the superstructure in one way or another, and focus on tree-updates as the primary means of accomplishing work within the applications. These modes may easily be combined within a single application.

- 5 Table 10 summarizes the various programming mechanisms, and their uses. These programming modes are next discussed.

Mode	Description	Benefits / Uses
SQScript	An XML-based programming language that offers the ability to write structured programs that modify the superstructure directly. SQScript programs may either make small attribute modifications to the superstructure, or may completely transform existing objects into new objects using arbitrary rules.	Completely client-local, this approach allows a SimpleOS application to be completely self-contained.
Network Server	In this mode, parts of the superstructure are converted into XML and sent to a remote device where a server-side programming language receives the data. Updates to objects are described on the network server and sent back to the device where they update the superstructure.	Allows for a program or part of a program to be controlled from a remote-server where additional data and other software resources can be found.
Client-Native	When supported, allows native code residing on the device to receive messages from the application and perform updates directly to the application's superstructure using a native-language API.	Allows existing programming logic to control a SimpleOS application, possibly with increased speed or access to client resources. This mode may be extended to the point that SimpleOS becomes an exclusive a shell for the application's user interface.
Default Handlers	Built-in event handlers that come with SimpleOS and perform common operations without any additional programming. .	May be used to perform simple actions without the need for any custom programming.

Table 10 - The Four Core Programming Modes

SQScript

- 10 SQScript, a further aspect of the present invention, is a simple and powerful client-local programming language useful in implementing various inventive features described herein . SQScript allows programmers to write methods on objects within the superstructure, or to create interpolated expression blocks simply by enclosing code inside {} blocks. See the "SQScript" discussion below.

15 **Default Handlers**

One of the four programming modes available in SimpleOS is the default handler. As a convenience to the SimpleOS developer, the environment provides a set of default

handlers that can perform simple operations such as changing the current screen, updating a simple list of attributes, pausing an application, etc. These default handlers are implemented internally using native code and affect the application's superstructure through an API within the code of the SimpleOS implementation itself.

- 5 The context argument passed to the default handler will have varying meanings depending on the nature and specification of the handler. Arguments to default handlers may be provided in the message descriptor.

Client-Native Code Handler

- 10 Some implementations of SimpleOS will support message handlers written in native code. These handlers comprise one of the four programming modes available from within SimpleOS. In this case, information inside of the message descriptor will indicate runtime dispatch information. The client-local runtime handler will have access to the superstructure from an API provided by SimpleOS.

- 15 The client-native API basically allows code to register for callbacks when particular events occur. Once control has been given to a client function, it can read and write values out of the superstructure, instantiate objects and modify data trees by using a set of API calls. The changes that the client-handler makes are verified the same way as any other change made by an application handler. See "Execution Cycle" below. Changes must conform to the basic superstructure rules, and there is no way to
- 20 perform any operation within the superstructure that would not also be allowed by SQScript.

- 25 In the most extreme case of using client-native code, SimpleOS becomes a user-interface shell around purely native application logic. In this sense, SimpleOS can be seen as a user interface library. A 3rd-party developer can link this user interface library into an executable program to give it the capabilities of SQUI. This product offers a large advantage for developers because it will save them the effort of writing custom GUI code for their own applications.

Remote Execution Model

The target of a message may be remote resource available over the Internet or on a wireless local area network such as a Bluetooth-enabled vending machine or kiosk. In this case, the message object is formatted and serialized into SQML (an XML grammar) that is sent to the target using the necessary protocol. In this case, a separate communication mechanism is used for the remote system to update the superstructure on the device. See "Network Communication" below.

Execution Cycle

SimpleOS (or any superstructure-based application environment in accordance with the invention) uses a message-driven model to deliver events and other information to the application. Each running SimpleOS process maintains a message queue with events from the outside world such as user activity, interface manipulations, network messages and external interfaces. Through data in the superstructure, the application describes which events it is interested in receiving. When the queue has one or more relevant messages, SimpleOS will dequeue each message one-by-one and deliver it to the application's requested handler. At any time, zero or more events may be waiting for a particular application, and they are processed one by one. This message queue can be a first-in first-out queue, with a mechanism to defer the delivery of messages if the application is unavailable or if network connectivity is necessary.

When an application's handler is invoked, it can receive information about the current state of the user interface and other external interfaces by reading data within its superstructure. The handler may then make modifications to the superstructure based on this data. When the handler exits, the operating system examines the changes to the superstructure and performs the relevant activities within the external systems that pertain to the modified sections of the superstructure. The model causes all changes to be atomic with respect to the superstructure, where each change is the result of a completed evaluation pass of the application code. The application never retains execution flow with its own event loop or polling system under any circumstances.

The execution cycle is comprised of the sequence shown in Table 11. Each of the five phases shown are next discussed.

- 5 1. **External Event Occurs:** Some event outside the application occurs, such as user activity or a network event. This is commonly the result of interaction with the user interface, such as a button press.
2. **Primary Update to Superstructure:** The operating system updates the superstructure if necessary with data from the event. For example, the new value of a text field might be written into the `value` attribute of a text field element.
- 10 3. **Invoking Application Handlers:** A registered application handler is selected and run. This handler may be any of the four types.
4. **Modifications to the Superstructure:** The handler modifies the superstructure in some way and exits. For instance, the application could change the value of a text label.
- 15 5. **Verification and Update:** SimpleOS verifies that the new state of the superstructure is correct and performs the modification. For instance, the application might see that a text label has changed in the superstructure and draw the change on the screen. SimpleOS re-renders the interface if necessary, potentially running interpolated code within the GUI definitions.

Table 11 - Phases of the Execution Cycle

20 **External Events**

A process may receive virtually any type of external event. These external events are typically user interface events such as button presses or field updates. However, they may also be network events, messages from other application processes or telephony activity. Bluetooth or 802.11b messages may also be processed inside an application.

- 25 In one practice of the invention, the context for an event is kept within an Event Object, which is instantiated into the superstructure with a temporary name and passed as an argument to the handler. Within the event object, there is data that may include the time and date of the event, details about how the event occurred and other event-specific data.

30 **Primary Update to Superstructure**

- Before an application handler is executed, SimpleOS must synchronize external state with data inside the superstructure. Depending upon the selected implementation, this may either occur all at once, just before the delivery of a message, or it can occur incrementally as the changes take place. Regardless, the application's code is not
- 35 executed unless all of the data inside the superstructure accurately reflects the state of the application.

Dynamic text fields present a perfect example of why this phase is important. When a user updates a text field, the application is not typically interested in knowing about each and every key press. Instead, the operating system will update the screen with each key press. However, when the user presses a “continue” button, the application probably will register an event handler so that it may verify the data within the text field and perform some other action. In this case, before the application’s code is run, SimpleOS must make sure that the superstructure has been updated with the value of the text field. It doesn’t matter if this occurs as the field is changing, or just before the application’s code is run.

10 Invoking Application Handlers

In one practice of the invention, the application can register an event handler using any of the four programming modes. A synopsis of the available modes is shown in Table 12.

Handler Mode	Description
Default Handler	A “canned” handler pre-defined in the system that performs a basic task such as changing the current screen.
SQScript	A local function stored in the superstructure. This handler is passed the name event object as an argument. See text for more information.
Client-Native Code Handler	A client-local handler written in some native language on the device. The code is passed the name of the event object.
Network Server	The event object is sent to a remote network location in the form of SQML. See text.

15 Table 12 - Synopsis of Event Handler Modes

Modifications to the Superstructure

When the handler runs, it may update sections of the superstructure using either a template expansion or a set of node updates. SimpleOS tracks these changes as they occur, so that the new state of the superstructure can be efficiently synchronized with the host operating system when the modifications are complete.

Verification and Update

Once the handler has completed, SimpleOS may verify the superstructure to ensure that it still contains valid information. This verification may include security checks,

validation of data formats and object references, and other steps that help make the application environment more secure.

Once SimpleOS is satisfied with the changes, it must perform the work that is described in the new version of the superstructure. For instance, if the application has
5 modified the labels of text fields, these text fields must be re-drawn. Often, this will involve copying data out of the superstructure and into the native programming libraries that SimpleOS uses within the host environment. This can also serve as stage where data can be verified for validity, security and consistency.

The change list compiled during the previous phase may be used at this point to
10 optimally make modifications. Since SimpleOS knows the sections of the superstructure that were changed, it can avoid performing activities within the host environment that are not required.

SimpleOS allows the application to insert dynamic SQScript expressions within the values of some SQML attributes and literal fields. Any attribute or literal field that
15 contains one or more SQScript expressions may change in value over the course of the program's execution. Whenever SimpleOS renders the current SQCard, it re-evaluates all SQScript expressions found in its attributes and literals and displays the new results. It is by means of these SQScript expressions that the programmer is able to cause computed data to be displayed. For instance, let's say the application is presenting a
20 set of check boxes to the user and wishes to display the number of checked boxes. The application could define a label for a text field that includes a SQScript expression block that computes the number of checked boxes. Whenever a change to the state of the application occurs, SimpleOS evaluates this expression block and displays the newly computed number of checked boxes. The programmer may also use this feature to
25 alter the appearance of the SQCard by inserting SQScript expressions into the SQML attributes that control style. Because this feature corresponds so closely to the **attribute value templates** defined by the XSLT specification, we use the same term to describe these inserted SQScript expressions in SQML.

Programming Contexts

In one practice of the invention, SimpleOS is configured to allow in-line code in SQScript to be executed in different places. This is not a requisite of the invention, but is useful in implementing the invention. Generally speaking, there are three occasions when
5 where programming code can be executed:

- **Handlers:** A SimpleOS application may register handlers for particular events on various objects. When the event occurs, the code is activated.
- **Attribute Value Templates:** Attribute value templates may appear within
10 some attribute values and literal text fields. They appear as SQScript expressions bounded by curly braces {}.
- **Initializer:** When SimpleOS instantiates a new object, it will run an *initializer* method on the object.

15 Each of these three code contexts applies selectively to the various programming modes available in SimpleOS. Table 13 describes which modes are available in which context. In cases where a mode is not available, an additional explanation is provided below.

Mode	Initializers	Event Handlers	Attribute Value Templates
SQScript	Yes.	Yes.	Yes.
Network-Resident	No. (Note 1)	Yes.	No, not even via SQScript. (Note 2)
Client-Native	No. (Note 1)	Yes.	No, not even via SQScript. (Note 2)
Default Handlers	No. (Note 1)	Yes.	No, not even via SQScript. (Note 2)

20 Table 13 - Availability of Modes within Code Contexts

Registering for Events

Those skilled in the art will appreciate that a superstructure-based application
25 environment must provide some means within its structural definition for an application to register handlers for events, so that when external events occur to the application and are updated by the operating system into the application's superstructure, the application has the opportunity to run an appropriate update, trigger some dynamic code, or take other action. In practice, these registrations could occur explicitly (as a

set of nodes that specify "when X occurs, run code Y"), or implicitly (through evaluation where the operating system attempts to evaluate some expression in the superstructure after an event has occurred and incidentally must force the evaluation of some code or the triggering of a network event). In either case, the registration could refer to local
5 superstructure code, client native code, code within the software container itself, or network events or communications.

Thus, in an example of a SimpleOS environment, event handlers must be registered inside the superstructure so that SimpleOS knows that an application is interested in a notification when an event occurs. Each registration must identify the
10 criteria of the event and provide a *target* where the event will be delivered. The following discussion describes two basic approaches for registering handler events inside SimpleOS.

In-Line Event Registration

The *in-line* style of event registration is the simplest and most convenient method for
15 registering for an event. In this style, elements within the data tree of an object contain an attribute named for a common event type, such as `onActivate` or `onChange`. When the corresponding event occurs, the value of the attribute is examined and executed.

In its simplest form, the in-line value may simply be a small SQScript expression block. In this case, the expression block will cause the current card to change to
20 `card_main`. See, for example, Table 14.

```
<sql:sq-card id="card_menu">
  .
  .
  .
  <sql:right label="CHECK OUT" class="border"
25      onActivate="{ui:card='card_main'}" />
  .
  .
  .
</sql:sq-card>
```

Table 14 - In-Line Event Registration with Code

30 The in-line registration value may also be a local method. In the example of Tables 15 and 16, pressing the "*Check Out*" button will invoke the method called next on the current object. The following example demonstrates this format.

```

5      <sql:sq-card id="card_menu">
        . . .
        <sql:right label="CHECK OUT" class="border"
          onActivate="next"/>
        . . .
      </sql:sq-card>

10    <function id="next">
      <sql:eval expr="index=$index+1" />
      <sql:if test="$index > $list_pizzaNames:size">
        <sql:eval expr="$index=1" />
      </sql:if>
    </function>

```

15 Table 15 - In-Line Event Registration for a Method

```

      <sql:sq-card id="card_menu">
        . . .
20    <sql:right label="CHECK OUT" class="border"
          onActivate="http://myserver.com/handleCheckout" />
        . . .
      </sql:sq-card>

```

25 Table 16 - In-Line Event Registration to a Server

The in-line registration may be a URL to a remote site. In this case, the current object is serialized and sent as a POST to the server shown. The server may respond to the POST with an update to the current object. The network communication model is described in detail below in the "Network Communication" section.

Finally, the in-line registration may refer to a message descriptor, as described in the section below.

Message Descriptors

Message descriptors allow for a more robust model for registering events. This mechanism is necessary since the in-line style of event registration does not provide a way to specify additional parameters for the receipt and dispatch of an event.

A message descriptor is an XML data type stored within the data tree of the appevent implicit object. SimpleOS will examine descriptors within the appevent object if no in-line message registrations apply. Each message descriptor defines the type of event, the element identifier for the entity that generates the event, and the location within the code section of the superstructure where the handler can be found. For instance, if the application wants to be notified when a button is pressed, it would

define a message handler of type "onActivate" for a source id of the button "myObject#myButton." The descriptor also must specify a code target, such as "anotherObject#root.code.myButtonHandler." This target pathname specifies the element identifier where the handler for the button can be found.

In addition to a type, source and target, the message descriptor may also contain a set of context elements that describe the ids of objects to be sent in the message. For instance, suppose the button "mainScreen#myButton" acts as a submit button to calculate data filled in on a form with fields "mainScreen#firstName" and "mainScreen#lastName". It would be helpful in the handler for myButton if the values of the buttons on mainScreen were easy to access. In this case, the context for myButton's handler could be set to the object "mainScreen", which would result in a copy of the mainScreen object and all of its data tree children to be easily accessible from within root.code.myButtonHandler. In cases where the message is delivered over the network, this syntax is imperative.

The chart shown in Table 17 summarizes the parts of a message descriptor.

Attribute	Description	Example
id	A unique id for the message identifier that is used to identify the message to the handler.	sampleEvent
Type	The type of event that will trigger this message	onActivate
Source	The source element for the event.	myScreen#myButton
Target	The message handler that will receive this message. In this case, a local method.	MyScreen#myButtonHandler
Context	An object to include in the event descriptor.	OtherScreen

Table 17 - Fields in a Message Descriptor

While the exact syntax for message descriptors may be left to the implementer, the hypothetical message descriptor definition shown in Table 18 may be employed. Thus, in Table 18, the descriptor will cause the handler myScreen#handler to be run when myScreen#myButton is activated. If the target was a network entity, the object otherScreen would be send along with the event object.

```
<sql:message-descriptor id="sampleEvent" type="onActivate"
  source="myScreen#myButton" target="myScreen#handler">
  <sql:context objectId="otherScreen"/>
</sql:message-descriptor>
```

Table 18 - Example: A Basic Message Descriptor

Message descriptors are commonly used to invoke native functions as well as default handlers.

Combination with In-Line Syntax

Message descriptors may be triggered from within the in-line syntax, or even by posting an event from within SQScript code. This is ideal for situations where a message must be sent over the network using the rich message descriptor functionality, but the programmer wishes to use the simplicity of the in-line syntax. For example, consider Table 19.

```
10 <sql:sq-class classId="class_main" objectId="main">
    <!-- The main card -->
    <sql:sq-card id="card_main">
      <sql:title label="MONOIDS" />
15   <sql:top label="ABOUT US" class="border" onActivate="#desc" />
      <sql:frame>
        <sql:image src="splash.gif"/>
      </sql:frame>
    </sql:sq-card>
20   <sql:message-descriptor id="desc" target="http://server.com/foo" >
      <sql:context objectId="otherScreen"/>
    </sql:message-descriptor>
25 </sql>
```

Table 19 - Example: Combining In-line Registrations with Message Descriptors

Here, when the onActivate event occurs on the top button, the message descriptor **desc**, stored on the current object, will be activated. This message descriptor will send a message object to **http://server.com/foo** that will include the value of the otherScreen object. Note that in this case the source and type attributes on the message-descriptor tag are not needed, since this context is provided by the in-line event registration.

Event Objects

When an event is delivered, an event object is temporarily created and passed as an argument to the message handler. The event object contains the context of the event, which may include the time that the event occurred, the sending object, the location of the cursor, and other information.

V: SQScript Language

The SQScript language, described by way of the following examples, discussion and reference materials, is a highly useful means to implement client-side code on SimpleOS. Unlike native code, SQScript can run on any of the SimpleOS supported
5 platforms. It provides stronger security than native-code handlers and is much easier to write. The following sections describe the SQScript language in detail.

It should be noted that the SQScript language demonstrated below is but one of many possible languages suitable for manipulating a superstructure-based application in accordance with the invention. A superstructure-driven operating system and
10 application environment could omit some features of the language described below, or add other features, and still be within the ambit of the invention. Other languages with different syntaxes could be used; or in certain cases, an implementer could avoid using any language locally at all, and instead rely on all updates to occur over a network. For purposes of the SimpleOS examples set forth herein, the applicants have opted to use a
15 local language.

It should also be noted that the programming language demonstrated below was designed to be syntactically similar to XSLT, a programming language used to transform tree structures. However, the execution environment whereby the application interacts with the operating system is (as will be seen below) radically different from XSLT,
20 because the tree transformations themselves form the basis of continuous processes of maintaining the application, as opposed to a one time translation.

Hello World Example

In connection with an introduction to SQScript, consider the following Hello World example, depicted graphically in FIG. 13. In this example, a screen 1300 is presented to
25 the user with a single button (such as a button of the conventional 4-way controller 1306, which includes arrow keys 1308, center key 1310 and additional keys below, 1312). When the button is pressed, the message "Hello World" is displayed. Note that a simple Hello World can actually be accomplished simply by defining a screen with the text "Hello World." However, this example involves some code, so it provides an

introduction to SQScript programming. In particular, possible source code for this application is shown in Table 20.

```

5  <sqlml xmlns:sqlml="http://www.sqgo.com/sqlml/sqlml.dtd"
    appid="com.sqgo.simpleos.demo.helloWorld"
    activate="main:card_main"
    >
10  <sqlml:sq-class classid="class_main" objectid="main">
    <!-- The main card -->
    <sqlml:sq-card id="card_main">
    <sqlml:title label="Hello World" />
15  <sqlml:right label="START" class="border"
        onActivate="{ui:card='main:card_hello'}" />
    <sqlml:frame>
    <sqlml:textarea>Welcome To the Hello World Application
    </sqlml:textarea>
20  <sqlml:textarea>Press "start" to continue
    </sqlml:textarea>
    </sqlml:frame>
    </sqlml:sq-card>
25  <!-- The "hello world" card -->
    <sqlml:sq-card id="card_hello">
    <sqlml:title label="Hello World" />
    <sqlml:frame>
    <sqlml:textarea>Hello World
30  </sqlml:textarea>
    </sqlml:frame>
    </sqlml:sq-card>
    </sqlml:sq-class>
35  </sqlml>

```

Table 20 - Example: Hello World

Note that the stylesheets for this and other application code examples have been omitted to improve the clarity of the examples. Bold sections of the code example are discussed below.

Within the initial sqlml tag, notice the presence of the attribute definition

activate="main:card_main". This tells SimpleOS that the initial card to show on the screen should be the entity named "main:card_main". In fact, SimpleOS will set the ui.card value to this string when the application loads. This statement presumes that somewhere within the process's superstructure, there is (or will be) an object called main with a sq-card definition called card_main.

Sure enough, in the sq-class definition, there is an attribute called `objectid`. This attribute indicates to SimpleOS that after the class `class_main` has been defined, an instance of the class will automatically be created and called `main`. This solves the problem earlier of having set the current card to `main:card_hello`.

5 Within the `right` tag, we define our one and only button. The purpose of this button is to change the current screen from `card_main` to `card_hello`. In order to do this, we have to register to receive the `onActivate` message on the right card button. We also must define some code so that the target of the handler will actually cause the screen to change. As previously discussed, screen changing is accomplished by
10 changing the value of `ui:card` to an element identifier that points to the desired object and card within the object.

The attribute `onActivate="{ui:card='main:card_hello'}"` performs both the registration and code definition action. (This is an example of an in-line event registration.) Let's take a closer look at this snippet. First of all, the expression
15 `ui:card='main:card_hello'` is surrounded by curly braces (the ones that look like this: `{}`). These curly braces tell the SQScript interpreter that when the `onActivate` expression is read, it should invoke the expression language parser to run the expression inside the braces. The expression plus its surrounding curly braces is called an *expression block*.

20 Finally, within the text area tag found in the second card definition, we find our "Hello World" text.

The SQML Tag

The `<sqml>` element type is the root of all SQML documents. The `<sqml>` element can contain an optional `<sqml:style-sheet>` element, followed by zero or more
25 `<sqml:sq-class>` elements. The `<sqml:style-sheet>` element is only allowed when the **doctype** attribute on the `<sqml>` element is set to **application** or **fragment**.

The formal XML declaration for the `<sqml>` element is as follows:

```
<!ELEMENT sqml ((style-sheet?), (sq-class*)) >
```

See Table 21.

Attribute	Purpose
xmlns	Required – defines the SQML namespace. It should be formatted as follows: xmlns:sqml="http://www.sqgo.com/sqml/sqml.dtd"
appid	Required – application ID. By convention uses a Java-like naming style. Example: appid="com.sqgo.simpleos.application.sample1"
sessionid	Required for applications with network interaction. Indicates the session to which this document applies.
doctype	Required, possible values include: application : this is a full SQML application document fragment : this is a fragment (update) to an SQML application message : this is a network message object : this is serialized object
activate	Names an object ID and sq-card to be activated by the application. Note that this is an object ID and not a class ID; therefore, the class should use the auto-instantiation feature and specify the same object ID. Example: <pre> <sqml xmlns:sqml="http://www.sqgo.com/sqml/sqml.dtd" appid="com.sqgo.simpleos.application.sample1" doctype="application" activate="main#card"> <sqml:sq-class classid="main"> <sqml:sq-card id="card"> ... sq-card elements ... </sqml:sq-card> </sqml:sq-class> </sqml> </pre>

Table 21 - The SQML Tag Attributes

5

Expression Blocks

The most basic execution unit within SQScript is an expression block. Expression blocks can be found almost everywhere inside a class or object. They indicate code in the SimpleOS Expression Language that should be evaluated. The SimpleOS Expression Language (EXL) is actually not based in XML like most of SimpleOS. The expression language is a string-based language with a special grammar vaguely reminiscent of XPATH, the expression language of XSLT. Expressions can involve assignment, variable lookups, data tree lookups, conditional expressions, mathematical calculations, predicate logic and a variety of other typical programming constructs. A substitution expression can be used to define the name or value of almost any attribute within the template as well as the predicate of any conditional or looping structure. The exact details of the grammar will be considered very shortly.

10

15

Expression blocks can be placed almost anywhere that the interpreter expects a string. When a string is needed (such as the value of the onActivate tag in the example above), the expression is immediately evaluated in its current context. When expressions are placed in an onActivate block, they are evaluated when the object in question receives an activate event. Expressions can also be found inside other places within the data tree, such as in the name of a label or the value of a text area. In this case, whenever SimpleOS needs to find the actual value of the string, it first will evaluate the expression and substitute the value of the expression into the place where it was found. In this case, the expression block is called a substitution expression. This behavior is also known as *expression interpolation*.

Consider the examples shown in Tables 22, 23 and FIG. 14. Each will result in a string "The magic number is 42." In these examples, the expression is evaluated every time the screen is redrawn.

```

15  <sql:sq-card id="card_hello">
      <sql:title label="Magic Numbers" />
      <sql:frame>
        <sql:textarea>The magic number is {2+40}.
      </sql:textarea>
20  </sql:frame>
    </sql:sq-card>

```

Table 22 - Example: Expression Interpolation in Text Elements

```

25  <sql:sq-card id="card_hello">
      <sql:title label="Magic Numbers" />
      <sql:frame>
        <sql:label text="The magic number is {2+40}." />
      </sql:frame>
30  </sql:sq-card>

```

Table 23 - Example: Expression Interpolation in Attributes

See FIG. 14 for visual presentation (screenshot) of the result.

Some attributes are not allowed to contain expression interpolation blocks. These include the "id" and "name" attributes, as well as within the structural tags such as sq-class.

Escaping Interpolation Characters

In order to use the { and } characters within a string without marking an expression block, you must escape them using the backslash character.

Interpolation and Assignment

If the interpolated expression is an assignment, the substituted value will be null. For instance, consider the following string:

5 "This is a{foo = \$bar} test."

When evaluated, this expression will become

 This is a test.

10 *Data Types*

Although all data in a SimpleOS application originates from SQML, which is entirely character data, the SQScript evaluator recognizes and handles data types other than character. SQScript supports the following data types:

- 15 • **String:** a string of characters.
 • **List:** a linear list of character strings
 • **Boolean:** true or false
 • **Numeric:** an integer or floating point value
20 • **Element Identifier:** a reference to an object or a named element in an object,
 such as a SqCard or variable.

Any quoted string in a SQScript expression is treated as a literal String value. Single or double quotes may be used. An unquoted string in a SQScript expression is an element identifier. When a literal Boolean value is called for, the built-in functions true() and
25 false() provide values. (The quoted string 'true' is a String value, and an unquoted string is an element identifier.)

 In this case, the string is parsed as an element or object identifier. If the element doesn't exist and the reference is followed, a runtime error will occur.

 The SQScript evaluator expects the expression found in the test attribute of an if
30 or for element to evaluate as Boolean. The evaluator never coerces a Boolean value

from a non-Boolean. A Boolean may be formatted as a String, the result of which will always be "true" or "false".

Arithmetic operators expect numeric operands, and produce numeric results. When evaluating arithmetic operators, the expression evaluator coerces string operands into numeric values by scanning the strings. If any operand string does not represent a number, the special value NaN (not a number) is the result of the operation.

Introduction to Variables

Variables within SQScript can be used to temporarily store data between evaluations of single expressions. The most basic type of variable is called a *scalar*. A scalar can represent any single string, number, Boolean or object/element identifier. Variables have a name that is used to find the value or to assign a new value.

Variables are declared by including a variable declaration somewhere within the scope of an object. "Somewhere" is very broad. If the variable declaration occurs directly underneath a class or object, it is considered a "member variable". However, a variable can also be declared anywhere within a function (we'll get to that later), in which case the variable is considered to have a scope bounded by the nesting of the program block. These rules are very similar to variable scoping rules in C or Java.

Table 24 shows an object with a single variable.

```
<sql:sq-class classId="myClass" extends="theObject">  
  <sql:variable name="index">1</sql:variable>  
</sql:sq-class>
```

Table 24 - Example: An Object with a Single Variable

In this example the variable is called index. Code inside myClass can reference this variable using a variable syntax like \$index. The examples shown in Tables 25 and 26 demonstrate the use of this syntax. (As defined herein, scalars can contain both numbers and strings.) The result of these two cards, when displayed, is shown in FIG. 15 (see 1502, 1504).

```

5  <sql:sq-class classId="myClass" extends="theObject">
    <sql:variable name="myString">Hello World</sql:variable>
    <sql:sq-card id="card_hello">
      <sql:title label="Example A" />
      <sql:frame>
        <sql:textarea>
10     The Value of myFirstString is "{$myString}".
        </sql:textarea>
      </sql:frame>
    </sql:sq-card>
  </sql:sq-class>

```

Table 25 - Example: Object with a String Variable

```

20 <sql:sq-class classId="myClass" extends="theObject">
    <sql:variable name="myNumber">3.1415926</sql:variable>
    <sql:sq-card id="card_hello">
      <sql:title label="Example B" />
      <sql:frame>
        <sql:textarea>
25     The Value of myFirstString is "{$myNumber}".
        </sql:textarea>
      </sql:frame>
    </sql:sq-card>
  </sql:sq-class>

```

Table 26 - Example: Object with a Numeric Variable

Variable Assignment

The := operator can be used to assign new values to variables. Consider the examples of Table 27, where the \$counter variable is used to maintain a running tally. Each time the user presses "Increment", the counter is incremented.

```

40 <sql xmlns:sql="http://www.sqgo.com/sql/sql.dtd"
    appId="com.sqgo.simpleos.demo.helloWorld"
    activate="main:card_main">
  <sql:sq-class classId="class_main" objectId="main">
    <sql:variable name="counter">1</sql:variable>
45   <!-- The main card -->
    <sql:sq-card id="card_main">
      <sql:title label="Variable Incrementing" />
      <sql:right label="Increment" class="border"
50       onActivate="{counter := $counter + 1}" />
      <sql:frame>
        <sql:textarea class="bigtext">
          Press Increment to increase the value:
60       </sql:textarea>
        <sql:textarea class="miditalictext">
          Current Value is: {$counter}
        </sql:textarea>
      </sql:frame>
    </sql:sq-card>

```

```
</sql:sq-class>
```

```
</sql>
```

5 Table 27 - Example: Variable Incrementing

Visually, the output of this program will look something like what is shown in FIG. 16 (see 1602, 1604, 1606).

10 **Variable Assignment Nuances**

In the previous example, the incrementing expression {counter := \$counter + 1} illustrates an important difference between counter and \$counter within the expression. The \$ tag performs the SQScript equivalent of a de-reference, looking up the value of the expression instead of representing the expression itself. If the expression had
15 merely been {counter := counter + 1}, SQScript would indicate a runtime error because counter on the left side of the expression indicates a pointer to the variable, not a number that can be added to 1.

Variable Methods and Attributes

Variables support a set of methods and attributes. Methods can be called on each
20 variable just like a method call to an object or class, and attributes can be accessed using the usual instance variable syntax.

Globally Available Variable Attributes

All variable types support the attributes shown in Table 28.

Attribute	Purpose
access	public private protected (default: private) Access modifier – works similar to C++ and Java; this only has meaning for variables declared at the class level. This attribute will cause an error to be thrown if used with variable declarations in other scopes.
final	true false (default: false) If true, value(s) cannot be modified. For Node type variables, setting final to true means that attribute values cannot be modified.
size	Returns the number of items in the variable; for a scalar variable (<sql:variable>) the size is always one. To get the length of the String value of a scalar variable, use the length method.

25 Table 28 – Summary Globally Available Variable Attributes

Scalar Attributes

In one practice of the invention, the following attributes (Table 29) are supported by scalar variables.

Method name	Purpose
length	Returns the length of the String value contained in a scalar variable. If used as an lvalue, will truncate or pad the string (with spaces) to the specified length. Example: truncates "ice cream" to "ice" <sqlml:variable name="myvar">ice cream</sqlml:variable> <sqlml:eval expr="myvar#length := 3"/>
tokenize(sep)	Splits the string into a series of tokens using the separators. This should be used as an rvalue in an assignment where the lvalue is the name of a list variable. Anything in the list prior to the string token operation is destroyed. Example: <sqlml:variable name="mystring">every good boy deserves fudge</sqlml:variable> <sqlml:list name="mylist"/> <sqlml:eval expr="mylist := \$mystring#tokenize(' ')/>

Table 29 – Scalar Attributes

List Methods

List variables support the following methods (Table 30):

Method name	Purpose
Remove(i)	Removes the item at position I
removelast	Removes the last item
removefirst	Removes the first item
Addlast	Adds an item at the end of the list
Addfirst	Add an item to the head of the list
Clear	Removes all items from the list
Append	Append the items in the rvalue list to the lvalue list. Example: appends list2 to list1 list1#append(\$list2)

Table 30 – List Methods

This

In order to explicitly avoid name collisions between inner and instance scope variables of the same name, **this** can be specified as an object qualifier. A variable with the name **this** is automatically created as an instance member. The value of **this** is set to the object ID. A new value cannot be assigned, nor can any other variable with the name **this** be created, regardless of context; i.e. **this** is a reserved word.

The example of Table 31 demonstrates the use of **this**.

```

<sq-class classid="thisexample">
  <sqlml:variable name="myvar"/>
  <sqlml:function id="init" type="initializer">
5    <sqlml:param name="myvar" type="string"/>
    <sqlml:eval expr="this#myvar := $myvar"/>
  </sqlml:function>
</sq-class>

```

Table 31 - Example of "this"

Introduction to Lists

SimpleOS allows variables to include more than one value. In this case, the variable is called a list. A list can be declared using syntax like that shown in Table 32.

```

<sqlml:list id="list_pizzaImages">
  <sqlml:list-element>thincrust.gif</sqlml:list-element>
  <sqlml:list-element>thickcrust.gif</sqlml:list-element>
20  <sqlml:list-element>deepdish.gif</sqlml:list-element>
</sqlml:list>

```

Table 32 - Declaring List Variables

The elements of the list may be accessed using array-index syntax, such as:

```

```

The first item is always 1, the second 2, etc. If a list expression appears within an attribute value template, the template is replaced by the elements in the list, separated by spaces. For instance:

```
<text-label label="{ $list_pizzaImages }" />
```

would generate a label:

```
thincrust.gif thickcrust.gif deepdish.gif
```

The number of elements in a list can be returned by using the count() function.

{ c unt(\$list_pizzaImages) } evaluates to 3

There are several other built-in functions for working with lists in SQScript.

List Example

5 Using arrays, a simple pizza chooser application can be created. This application will let the user cycle between a list of pizzas on a card, seeing both a photo and a text description of each. Because we have not covered conditional expressions yet, there is nothing to prevent the user from cycling beyond the three pizzas we have defined.

(This example will be revisited later in the section on methods.)

10 In the source code shown in Table 33, the system maintains a variable called "index" that tracks the number of the current pizza.

```
<sql:sq-class classid="class_menubase" id="main">
15 <sql:variable name="index">2</sql:variable>

    <sql:list name="list_pizzaNames">
        <sql:list-element>Thin Crust Pizza</sql:list-element>
        <sql:list-element>Thick Crust Pizza</sql:list-element>
20 <sql:list-element>Deep Dish Pizza</sql:list-element>
    </sql:list>

    <sql:list name="list_pizzaImages">
        <sql:list-element>thincrust.gif</sql:list-element>
25 <sql:list-element>thickcrust.gif</sql:list-element>
        <sql:list-element>deepdish.gif</sql:list-element>
    </sql:list>

    <sql:sq-card id="card_menu">
30 <sql:title label="Pizza Menu" />
        <sql:top label="PREVIOUS" class="border"
            onActivate="{index=$index-1}" />
        <sql:bottom label="NEXT" class="border"
            onActivate="{index=$index+1}" />
35 <sql:frame>
            <sql:image src="{list_pizzaImages[$index]}" />
            <sql:textarea>
                {list_pizzaNames[$index]}
            </sql:textarea>
40 </sql:frame>
    </sql:sq-card>

</sql:sq-class>
```

45 Table 33 - Example: Pizza Ordering Version 1.0

The "next" and "previous" buttons have been wired up to increment and decrement the index variable. Within the frame definition, an image and a text area extract information

from two arrays list_pizzaNames and list_pizzaImages based on the current value of index.

Screenshots corresponding to this example are shown in FIG. 17 (1702, 1704).

Object Lists, Maps and Vectors

- 5 Although lists, hash maps and vectors can only store strings as the primitive object types, they can in fact be used as very powerful collection classes to organize groups of objects. Since all objects have a unique ID, which is a string, creating a list, hash map or vector of object IDs is equivalent to storing references to the objects.

10 By the same token, IDs of XML nodes could also be stored in these collection classes, and then the IDs could be used to dynamically build Node variables to access an object's DOM structure. Be aware, however, that the XML IDs are only guaranteed to be unique within an object instance. So be careful not to mix references to different object's DOMs.

15 When an object is instantiated using the **<sqlml:new>** statement, an automatic variable is created in the current scope called newobjectid. The example shown in Table 34 builds a list of objects that contain messages.

```
20 <sqlml:sq-class classid="class1">
  <sqlml:message id="msg1"/>
</sqlml:sq-class>
<sqlml:sq-class classid="class2">
  <sqlml:message id="msg2"/>
</sqlml:sq-class>
25 <sqlml:sq-class classid="class3">
  <sqlml:message id="msg3"/>
</sqlml:sq-class>
<sqlml:sq-class id="main" objectid="main">
  <sqlml:list name="objectList"/>
  <sqlml:function name="init" type="initializer">
30   <sqlml:new classid="class1"/>
   <sqlml:eval expr="objectList#addlast($newobjectid)"/>
   <sqlml:new classid="class2"/>
   <sqlml:eval expr="objectList#addlast($newobjectid)"/>
   <sqlml:new classid="class3"/>
35   <sqlml:eval expr="objectList#addlast($newobjectid)"/>
  </sqlml:function>
</sqlml:sq-class>
```

Table 34 - Example: Building a List of Objects

Implicit Tree Variables

Certain elements within the tree data of an object also create implicit variables. Most GUI elements use this approach to provide a way to extract their current value. For instance, a text field declaration automatically creates a variable with its id:

```
<sqlml:text-field id="username" label="Enter your name"
value="foo" />
```

The current contents of the text field can be extracted by referring to \$username.

Consider the code example shown in Table 35, which creates a text field and a text label. When the "GO" button is pressed, the value of the text field is copied into the current value.

```
<sqlml xmlns:sqlml="http://www.sqgo.com/sqlml/sqlml.dtd"
appid="com.sqgo.simpleos.demo.textFieldExample"
activate="main:card_main"
>

<sqlml:sq-class classid="class_main" objectid="main">
  <sqlml:variable name="currentValue"></sqlml:variable>

  <sqlml:sq-card id="card_main">
    <sqlml:title label="Text Field Example" />

    <sqlml:right label="GO" class="border"
      onActivate="{currentValue := $username}" />
    <sqlml:frame>
      <sqlml:textarea>Enter Text</sqlml:textarea>
      <sqlml:text-field id="username" label="Your Text"
        value="" />
      <sqlml:textarea>Current Value is: {$currentValue}
    </sqlml:textarea>
    </sqlml:frame>
  </sqlml:sq-card>
</sqlml:sq-class>

</sqlml>
```

Table 35 - Example: Text Field Extraction

This example will look something like that shown in the screenshot of FIG. 18 (1802, 1804), when executed. The user has already entered the string "Jeremy" into the text field.

Extended Expressions

Just like C and Java, SQScript expressions can contain basic mathematical operation. This may include arithmetic as well as mathematical functions. Function use the syntax `function()`, where a variable or set of variables is included within the parenthesis separated by commas.

Some examples of valid expressions with and without functions are shown in Table 36.

```
{ foo := 12/16 }  
{ foo := 12/(16+23+19+$bar) }  
{ foo := cos( $bar) }
```

Table 36 - Examples of Mathematical Expressions

Instance Variables

If a list or scalar variable is declared underneath the class or object definition, the variable becomes an instance variable. This allows it to be accessed from other objects using the standard element identifier syntax. In the example shown in Table 37, the object `main2` loads the values of `main:username` and `main:var`.

```
<sql:sq-class classid="class_main" objectid="main">  
  <sql:sq-card id="main">  
    <sql:title label="TEST" />  
    <sql:frame>  
      <sql:text-field id="userName" label="Type Your Name"  
        value="Jeremy Gilbert" />  
    </sql:frame>  
  </sql:sq-card>  
  
  <sql:variable name="myScalar">This is a scalar</sql:variable>  
</sql:sq-class>  
  
<sql:sq-class classid="subclass" objectid="main2">  
  <sql:sq-card id="main2">  
    <sql:title label="TEST TWO" />  
    <sql:top label="ABOUT US" class="border"  
onActivate="{ui:card='main#card_about'}" />  
    <sql:frame>  
      <sql:text-label id="userName" label="Your Name Is"  
        value="{ $main:userName } and { $main:var }" />  
    </sql:frame>  
  </sql:sq-card>  
  
</sql:sq-class>
```

Table 37 - Example: Instance Variables

Methods

SQScript methods are almost identical to methods in other object-oriented languages. Each method runs inside the variable binding context of the object it was invoked upon.

- 5 That means that local variables and data within the object's data tree are locally accessible to the code of the method. Methods can be invoked within any SQScript expression, including from other methods.

Defining Methods

- 10 Methods may be defined using the `sqlml:function` tag. Referring now to Table 38, the example object `main` of class `class_menubase` defines a method called `previous` that decreases the value of the variable `index` by 1.

```
15 <sqlml:sq-class classid="class_menubase" id="main">
    <sqlml:variable name="index">2</sqlml:variable>

    <function id="previous">
20     <sqlml:eval expr="{index := $index-1}" />
    </function>

</sqlml:sq-class>
```

25 Table 38 - Example: Method Definition

The Method Body

- 30 Within the body of the method, statements can be included in an ordered list. These statements may include conditional statements, expressions, loops and other basic programming structures.

When a method is invoked, the statement elements underneath the definition are evaluated one-by-one. When the last statement is evaluated, the method terminates.

- The most common type of statement within a method body is an `eval` statement.
- 35 The `eval` statement element takes a single attribute `expr` that contains a SQScript

expression. The curly braces are required here and every other time a SQScript expression is included somewhere within the superstructure.

The example of Table 39 demonstrates multiple expressions chained together. At the conclusion of the method, the value of \$test will be 1111.

```

5  <sql:sq-class classid="base_class" id="main">
    <sql:variable name="test">1</sql:variable>
10  <function id="doStuff">
    <sql:eval expr="{test=$test+10}" />
    <sql:eval expr="{test=$test+100}" />
    <sql:eval expr="{test=$test+1000}" />
15  </function>
    </sql:sq-class>

```

Table 39 - Example: Evaluation Tags in a Method

Invoking Methods

Methods may be invoked from within any SQScript expression. The syntax for method invocation is the name of the method followed arguments bracketed within a set of parentheses. Even in cases where there are no arguments, the parentheses are needed to denote a method invocation to the parser.

In the example of Table 40, the method doStuff invokes a second method called helpMe.

```

30  <sql:sq-class classid="base_class" id="main">
    <sql:variable name="test">1</sql:variable>
    <function id="doStuff">
    <sql:eval expr="{test=$test+10}" />
    <sql:eval expr="{:helpMe()}" />
35  </function>
    <function id="helpMe">
    <sql:eval expr="{test=$test+100}" />
    <sql:eval expr="{test=$test+1000}" />
40  </function></sql:sq-class>
    </sql:sq-class>

```

Table 40 - Example: Method Invocation

In addition, methods may be called using a variety of formats, shown in Table 41.

5	<code>main:helpMe()</code>	Indicates a call to method <code>helpMe</code> on <code>main</code>
	<code>:helpMe()</code>	Calls method <code>helpMe</code> on the current object
	<code>helpMe()</code>	Calls method <code>helpMe</code> on the current object

Table 41 - Method Invocation Examples

10 **Variables Scope**

Scalar and list variables may be declared inside methods. These variables are not visible outside of the method under any circumstances, and are only available to code within the method. If a variable within a method (or a block) has the same name as a variable outside of the block or a method, the variable will obscure the definition further away from the currently executing code.

When the code shown in Table 42 runs, the inner copy of `test` is incremented, not the outer copy of `test`. However, it is still possible to reach outer instances of variables by using element syntax. The code shown in Table 43 increments both instances of `$test`.

```
20 <sqlml:variable name="test">0</sqlml:variable>
    <function id="doStuff">
25   <sqlml:variable name="test">0</sqlml:variable>
      <sqlml:eval expr="{test := $test + 10}" />
    </function>
```

Table 42 - Variable Masking

```
30 <sqlml:sq-class classid="base_class" id="main">
    <sqlml:variable name="test">0</sqlml:variable>
35   <function id="doStuff">
      <sqlml:variable name="test">0</sqlml:variable>
      <sqlml:eval expr="{test := $test + 10}" />
      <sqlml:eval expr="{main:test := $main:test + 10}" />
40   </function>
</sqlml:sq-class>
```

Table 43 - Variable Masking with Element Identifiers

45 **Declaration Order**

The order of the variable declarations is not critical. The portions of code shown in Table 44 are equivalent and both valid.

```

<function id="doStuff">
  <sql:variable name="test">0</sql:variable>
  <sql:eval expr="{test=$test+10}" />
</function>

<function id="doStuff">
  <sql:eval expr="{test=$test+10}" />
  <sql:variable name="test">0</sql:variable>
</function>

```

Table 44 - Examples of Declaration Order

This results from the fact that variable declarations are not statements and are not parsed at the time that the statements are evaluated. Before a line of code inside a block is ever interpreted, the variables inside are collected together to provide a context object for the evaluator. The context object does not care where the definition occurs, only where it is within the block structure.

Blocks and Variable Scope

Most looping and conditional statements create new variable context where it is possible to mask variable names. Within a method, variables will bind to their closest block. Consider the example shown in Table 45. Here, the assignment `{index := 1}` within the if block will only update the inner index variable.

```

<function id="next">
  <sql:variable name="index">0</sql:variable>          # Not updated
  <sql:eval expr="{index := $index+1}" />
  <sql:if test="{index > 5}">
    <sql:variable name="index">0</sql:variable> # Updated
    <sql:eval expr="{index := 1}" />
  </sql:if>
</function>

```

Table 45 - Blocks and Variable Scope

Conditional Statements

One of the most basic control structures is the if block. The if block only evaluates expressions inside of the block if the test attribute evaluates to true. The if block begins by running the code within the test expression. If the result is true, the evaluator is allowed to continue into the block. Otherwise, the block is skipped. Conditional statements only work inside methods.

In the example shown in Table 46, the variable numItems is checked. If the value is greater than 100, the variable \$message is updated.

```

5      <sqlml:variable name="numItems">0</sqlml:variable>
      <function id="checkLabel">
        <sqlml:if test="{ $numItems > 100 }">
          <sqlml:eval expr="{message := 'You have too many items' }" />
        </sqlml:if>
10     </function>

```

Table 46 - Example: If Block

Loops

15 SQScript can also support looping structures.

Pizza Ordering Revisited

The example shown in the screenshots of FIG. 19 (1902, 1904, 1906) demonstrates a number of concepts. The application allows the user to cycle through a set of three pizza selections. However, unlike the previous pizza ordering example, this version
 20 maintains a check to ensure that if the user goes past the last item in the list, the selector will wrap back to the first item.

The example of code shown in Table 47 operates by defining two methods, next and previous, that increment or decrement a counter with the current pizza name. They also contain a check to reset the value if needed.

```

25
      <sqlml:sq-class classid="class_menubase" id="main">
        <sqlml:variable name="index">2</sqlml:variable>
30     <sqlml:list name="list_pizzaNames">
      <sqlml:list-element>thin crust pizza</sqlml:list-element>
      <sqlml:list-element>thick crust pizza</sqlml:list-element>
      <sqlml:list-element>deep dish pizza</sqlml:list-element>
35     </sqlml:list>

      <sqlml:list name="list_pizzaImages">
      <sqlml:list-element>thincrust.gif</sqlml:list-element>
      <sqlml:list-element>thickcrust.gif</sqlml:list-element>
40     <sqlml:list-element>deepdish.gif</sqlml:list-element>
      </sqlml:list>

      <sqlml:sq-card id="card_menu">
      <sqlml:title label="Pizza Menu" />
45     <sqlml:top label="PREVIOUS" class="border"
        onActivate="{previous()}" />
      <sqlml:bottom label="NEXT" class="border"
        onActivate="{next()}" />

```

```

5      <sqlml:frame>
      <sqlml:image src="{list_pizzaImages[$index]}" />
      <sqlml:textarea>
        {list_pizzaNames[$index]}
      </sqlml:textarea>
    </sqlml:frame>
  </sqlml:sq-card>

10  <sqlml:variable name="index">1</sqlml:variable>

    <function id="next">
      <sqlml:eval expr="index := $index+1" />
      <sqlml:if test="$index > $list_pizzaNames:size">
        <sqlml:eval expr="$index := 1" />
15    </sqlml:if>
    </function>

    <function id="previous">
      <sqlml:eval expr="index := $index-1" />
      <sqlml:if test="$index = 0">
        <sqlml:eval expr="{index := $list_pizzaNames:size}" />
20    </sqlml:if>
    </function>

25  </sqlml:sq-class>

```

Table 47 – Example: Pizza Ordering Revisited

Inheritance

30 SQScript supports the object-oriented programming concept of inheritance between classes. Each object that has been instantiated in the system belongs to a class. When methods and variable lookups occur on an object, the definition for the lookup is found by searching the inheritance tree starting at the object. First the class is searched, then the super-classes in order.

35 When an object is cloned from another object, the new object assumes the class of the original.

Inheritance does not serve any typing ability. Classes exist as organizational means for grouping methods and data members.

Creating Objects

40 Object manipulation is an important aspect of the SQScript language, since objects serve as the basic grouping for data within the superstructure. When the process is first loaded, there are no instances of any objects except the implicit objects. Each class is loaded into memory, and one of them is used to create the initial object. A programmer

must declare at least one class within the SQML source to have an objectID. This causes an instance of the class to be automatically generated at runtime.

After the SimpleOS program starts, the developer may create additional object instances. Objects are created from templates, which can be either another object or a class. If the template for a new object is an existing object, the new object can become a carbon copy of the old object. The new object will assume the same class affiliation as the old object. If the template for a new object is a class, all of the data values within the class are copied over into the new object.

No two objects in the process's superstructure may have the same instance name. If a new object is created with a pre-existing instance name, the old object is completely replaced by the new object. In some cases, a programmer may wish to create a temporary object and want to ensure that the name does not conflict with a previously created object. A special programming primitive allows the programmer to create a new name for an object that is guaranteed to be unique.

Syntax

The basic object operation is "instantiate". This operation takes an objectid attribute that describes the new instance name for an object. The source attribute can be used to specify either the class or object id that forms the template for the new object.

```
<function id="makeObject">  
  <sql:instantiate objectid="options" source="class_options" />  
</function>
```

Initializers

When an object is first instantiated, a method on it called initialize is called. This gives the object an opportunity to perform any initial configuration.

VI. Superstructure Operations

Object Transformations

SimpleOS provides an advanced means for object manipulation called *object transformation*. Object transformation is similar to the normal process of object instantiation, however an additional processing pass is taken before the object is created. A transformation is a complex expression organized as a dynamic template that maps old data values onto new data values. The template contains embedded expressions that allow nodes to be reorganized, repeated, manipulated and altered.

During an object transformation, a template (represented as a tree) is loaded from the superstructure. This template describes a functional mapping between an existing object and a new object. When SimpleOS evaluates a SQScript template during an object transformation, various instructions within the template cause substitutions within this template. Substitutions can include finding other data within the tree, looping and indexing, conditional expressions and other common programming structures. Once the template has been filled out, the resulting value becomes the definition for the new object.

Partial Transformations

For efficiency, SimpleOS also allows a transformation to express a “clipped” tree instead of the complete definition of a new object. In the clipped tree model, the result of a template transformation is a partial list of updates that are selectively folded back into the object. A clipped tree is not a sub-tree. It faithfully follows the root-to-node structure and naming of the elements within the object. However, the clipped tree can be sparse – it does not have to include unaffected siblings or sub-trees. Every node in the clipped tree must have an expressed path to the root, but only the nodes that must be updated, changed or deleted need to be in the clipped tree itself.

One advantage of partial transformations is the ability to track the state when the clipped tree is applied to the object. The folding process tallies all of state changes as the tree is updated. This provides rich information that SimpleOS uses to create a

minimal set of work required to perform changes to the environment. This list of changes can be used for optimal GUI updating and for applying rules about how the user interface migrates from one state to another through animation, etc.

Messaging

5 Messaging is a special type of instantiation of a class over a network. The term messaging, as used in this concept, is entirely unrelated to Message Queues or event messages used by SimpleOS. In this context, messaging refers only to network communications. The target of a message can be a device or a server. Therefore, it is necessary for all implementations of SimpleOS and SimpleServer to provide the
10 SQScript environment.

A message is an SQML document containing one or more SQClass templates. The type attribute for the **<sql>** element is **message**. SimpleOS generates the required **appid** and **sessionid** attributes. In addition, SimpleOS will also provide a certification digest that is not documented here for security reasons.

15 Table 48 shows an example of a message.

```
20 <sql xmlns:sql="http://www.sqgo.com/sql/sql.dtd"
    appid="com.sqgo.simpleos.application.sample1"
    sessionid="8A7B6C5D4E3F8A7B6C5D4E3F8A7B6C5D"
    doctype="message">
    <sql:sq-class classid="class_main">
      <sql:message id="msg">
        Some random text data
      </sql:message>
    </sql:sq-class>
25 </sql>
```

Table 48 - Example: Body of a Message

30 A number of factors are to be considered with regard to messages:

- The objects in the message are instantiated on the target, not the source.
- All objects are automatically instantiated for all classes defined in the message
35 body and any initializer functions are executed.
- Other functions within the instantiated objects may be invoked by the target scripting environment.
- Object IDs are automatically assigned by the target system – any values supplied for the **<sql:sq-class>** or **<sql:new>** **objectid** attribute are ignored. The

reason this is necessary is that the source (which emits the SQLML) cannot be aware of the object namespace that exists on the target.

- The message is actually serialized from the application's internal DOM structures into the text representation of SQLML prior to sending over the network. The transport protocol is assumed to be a request/response type of protocol like HTTP/HTTPS. The response may be a simple acknowledgment, or it may be another SQLML formatted message, depending on the needs of the application.

It is recommended that messaging occur over a secure connection whenever possible,

and that the connection be protected by a digital certificate of verifiable origin.

Intrinsic Methods

In the Java programming languages, all objects are ultimately inherited from the base class called *Object*. The Object class carries with it certain methods that are therefore available to all Java objects. It is not necessary for any class to specifically name Object as a superclass – it is implied in the language.

In SimpleOS, objects also inherit from an implied base class, although this class does not have a name. Methods from this class are available to any class defined in SQScript. The intrinsic methods are shown in Table 49.

Method name	Purpose
getobjectid	Returns the object ID associated with this object instance.
getclassid	Returns the class ID associated with this object instance.
getsuperid	Returns the class ID of the super class, if any.
getinterfaces	Returns a list of interfaces implemented by this object. Should be used as an rvalue, with a list variable as the lvalue. Any items existing in the list prior to the call will be destroyed.
save	Serialize this object to its XML representation
load	Restore this object from its serialized form; returns the new object ID.

Table 49 - Summary of Intrinsic Methods

Object Serialization

Object serialization in SimpleOS is designed to rely upon XML. Ultimately, all SimpleOS objects will be serializable in XML form, although this feature is not implemented in the reference implementation and may not be for several releases.

Ideally, serialized applications, and even processes, could be moved from device to device without interrupting operation.

Two intrinsic methods are provided for object serialization:

1. save – automatically creates message objects within the DOM that carry all internal object data, and then writes the DOM as XML text.
2. load – reads serialized XML text and restores the internal data structures.

5

A special version of the <sqml> element provides the wrapper for serialized objects. Multiple objects can be saved in one SQML stream. The example of Table 50 shows a serial data stream representing objects and classes inside a message.

```

10 <sqml xmlns:sqml="http://www.sqgo.com/sqml/sqml.dtd"
    appid="com.sqgo.simpleos.application.sample1"
    sessionid="8A7B6C5D4E3F8A7B6C5D4E3F8A7B6C5D"
    doctype="object">
15   <sqml:sq-class classid="class1">
    <sqml:message id="sqgocontext_123A4F9">
    ... Actual data not shown ...
    </sqml:message>
    </sqml:sq-class>
20 </sqml>

```

Table 50 - Example: Serialized Data

VII: Network Communication and Remote Execution

The following discussion describes ways in which SimpleOS can be used to support network communication and remote execution. Again, it is intended to illustrate and provide examples of implementation of the broader concepts of the invention, particularly as they relate to superstructure-based application environments. Accordingly, while the following discussion is based on SimpleOS and related examples, it will be understood that the principles illustrated can be implemented in a wide range of different configurations, all within the ambit of the present invention.

SimpleOS provides a novel model for network communication between the client and server if an application calls for such interaction. The developer is free to write server-side logic in any XML-enabled server programming language including Java, C++, Perl, Tcl, Shell and many others. Using an innovative methodology, small pieces of the application super-structure itself are shifted back and forth between the device and the server, allowing virtually any operation on the device to be performed by the server. This complete flexibility allows the programmer to adopt any request/response paradigm he chooses. Because of the abstraction offered by the application

superstructure, the developer has continuous flexibility to locate application logic either on the device or on a remote server.

Simple Quick GO transcends the existing client server and HTTP-based computing architectures. In client/server the primary exchange is data in some private or public form. In HTTP, the exchange is HTML and form information that conforms to HTTP. However, SimpleOS exchanges actual pieces of the application, large or small. Messages from the server **incrementally** update the application superstructure through the same folding process used by object transformations.

When a running application needs to communicate with a server, pre-selected objects from the process's superstructure are automatically sent to the remote site along with the event object. The remote server's code has access to any data that it needs since the same application superstructure available on the device has been transferred over the network. If the server needs to produce some change on the device, it can send a minimal set of changes back to the device along the same communication channel. These changes are applied to objects in the superstructure in the same way that code running on the device can update the superstructure using the transformation model. This entire process works automatically since it is hooked into the transformation engine used for object transformations.

In the remote execution model, messages normally delivered to the client's code are shipped to a remote server. In client-side operation, user events typically cause a message descriptor to be identified and filled out to produce a message object, which is handled as an argument to a template-based transformation. In the remote model, this message object is serialized into SQML (or any suitable, equivalent transport mechanism) that is sent to the target using any of the supported protocols. The server reconstitutes the device's message object from the SQML and performs its processing. Because SimpleOS includes any arbitrary copy of objects inside the application's superstructure, the server has access to all of the information it needs to process a request.

After processing a request, most applications typically call for an update of the application to confirm the action or refresh a display. SimpleOS handles the return

update by allowing the server to construct a result-tree object that minimally describes the necessary changes to any object within the superstructure. This result-tree object is the same data structure produced via template evaluation on the device. The server converts this response into SQML and sends it back to the device where it is un-

5 marshaled and applied to the superstructure. Because of the parity between the client and remote execution model, any activity possible in client code is also possible on the server.

Server Components

In one practice of the invention, to assist with secure remote application processing,

10 Simple Quick GO will provide a J2EE-compliant container that performs the work of translating messages from devices into a common form and then transmits a response back from program code. The present implementation, available first on the Apache/Tomcat platform, should work on any J2EE-certified application server platform. Developers can write JSP files or Java classes that can receive messages and perform

15 updates on applications running under SimpleOS.

In one practice of the invention, the SimpleQuickGO Server platform maintains a state-full session bean for each active SimpleOS application that has started a networking session. Every time a SimpleOS client application sends a message to the server, the software locates the individual session bean associated with the application.

20 This session bean receives all communication from the client and delegates messages to application-specific code. By extending Simple Quick GO's session bean interface, the developer can build any type of server-side functionality. The server platform is also responsible for securing the communication. (See Security section.)

In a further practice of the invention, the application server maintains a set of

25 JSP utilities and tag libraries to simplify the creation of SQML-based responses to the device. These utilities help parse messages from the device and construct the correct responses.

The application server provides an extensive set of services to aid in the provisioning of applications. The server can maintain a registration of known programs

and known terminals, and can field provisioning requests from devices that wish to load a new program. The server then will help choose and merge the stylesheets that apply to the particular device and create the provisioned application object. In some cases, media assets (such as sounds, images and animations) require additional translation,
5 for instance an image may need to be resized on the server before it is sent to the device.

Other Options

The SimpleOS network communication model can also work without the J2EE server component. Since the message transfer format is entirely XML-based, any programming
10 language capable of receiving XML over the network and emitting XML in response should be able to communicate with a SimpleOS application.

Updating Application Segments and Resources

SimpleOS applications can be dynamically updated. Portions of the application may be added, removed or replaced. However, changes are made in the application object, so
15 they do not become apparent until a new process is activated.

The mechanism for modifying an application involves sending a particular type of SQML document to the device, known as a *fragment*. A fragment contains portions of the application that will be changed. If the element structure defined by the fragment is new, it is added to the application. If the element structure already exists, the old
20 element structure is first deleted and then replaced with the new structure supplied in the fragment. If the element structure in the fragment is empty, then the corresponding element structure is removed from the application.

A fragment may consist of one or more sq-class definitions, and/or a style sheet. Individual items within classes or style sheets cannot be updated. Developers should be
25 careful not to create fragments that damage the integrity of the application's structure. While some checks are made (for example, not allowing referenced superclasses to be deleted), complete verification by SimpleOS is impractical.

Table 51 shows a fragment that would add or update a simple class in an existing application.

```

5  <sql xmlns:sql="http://www.sqgo.com/sql/sql.dtd"
    appid="com.sqgo.simpleos.application.sample1"
    doctype="fragment">
    <sql:sq-class classid="class1">
    <sql:list name="mylist"/>
10  <function name="init" type="initializer">
    <sql:eval expr="mylist#addlast('abc')"/>
    <sql:eval expr="mylist#addlast('it is easy as')"/>
    <sql:eval expr="mylist#addlast('123')"/>
    </function>
    </sql:sq-class>
15 </sql>

```

Table 51 - Using Fragments to Update an Application

Client/Server Example

20 Fragments can be used as the basis of an automatic style of communication between an application and a remote server. The following example demonstrates how the client device and the server might work in a fictitious example called NewsSubscribe. This application queries a user for a name and email address and then submits the data to a remote server. The application is very simple, and all of the communication in this
25 example is automatic. Aside from the initial setup described below, no additional set-up is required.

The following is merely a hypothetical example, in which the grammar utilized does not necessarily correspond to that described above (thereby illustrating that various grammars can be utilized in accordance with the invention). Note also that the
30 following example does not show the object-level grouping described above. The example does, however, demonstrate a method of using fragments to accomplish the desired result. The reader should also note the method of using objects to implement this function, as described above in the discussions regarding messaging. Other examples are, of course, possible and consistent with the invention. In this example,
35 the following nodes describe the application (see Table 52):

Node Name	Description	Attributes
news.gui	The GUI handler, set to use screen1 as the current screen.	currentScreen=screen1
news.gui.screen1.fullName	The GUI label to receive the full name	type=textfield value="" prompt="Full Name:"
news.gui.screen1.email	The GUI label to receive the email address	type=textfield value="" prompt="Email Address:"
news.gui.screen1.submit	The GUI button to submit the data	type=submit prompt="Submit"
news.mdsc.submitDesc	The message descriptor	Source=news.gui.screen1.submit Type=onActivate Target=http://myserver.com/sys Context=news.gui.screen1
news.gui.screen2.thankYou	A GUI label on a "thank you" screen	

Table 52 - Nodes

5 The interaction could occur as follows:

- 1) The user downloads or starts up the Subscribe SimpleOS application
- 2) The application loads and presents a screen to the user with a prompt for a full name and email address
- 10 3) The user enters this information and presses "submit"
- 4) The system consults news.mdsc.submitDesc and forms a new message object:

```

<sqml>
<message id="news.temp1.message" caller=" news.mdsc.submitDesc">
15   <context path="news.gui">
       <screen id="screen1">
           <textfield id="fullName" value="Jeremy Gilbert"/>
           <textfield id="email" value="jeremy@sqgo.com"/>
           <button id="submit"/>
20       </screen>
       </context>
</message>
</sqml>

```

- 25 5) The system checks for an existing connection to http://myserver.com/sys. Finding none, it establishes an HTTP connection to the server and provides an handshake message to the server such as:

```

<sqml>
30   <network-control seq=0 type="http" action="begin"
       terminalid="A23BE3FEC23223A"/>
</sqml>

```

- 6) The remote server sets up a new session handler for the "news" application, and returns the following message:

```
5  <sqml>
    <network-control seq=1 type="http" action="begin-connect"
      terminalid="A23BE3FEC23223A"
      sessionid="B44023957FA992E4EE"
      application="news" />
10 </sqml>
```

- 7) Now the application sends the message object it formed earlier, with network control information.

```
15 <sqml>
    <network-control seq=2 type="http" action="handleEvent"
      terminalid="A23BE3FEC23223A"
      sessionid="B44023957FA992E4EE"
      application="news" />
20
    <message id="news.temp1.message" caller=" news.mdsc.submitDesc">
      <context path="news.gui">
        <screen id="screen1">
          <textfield id="fullName"
25             value="Jeremy Gilbert"/>
          <textfield id="email" value="jeremy@sqgo.com"/>
          <button id="submit"/>
        </screen>
      </context>
30 </message>
</sqml>
```

- 8) The server passes a parsed XML DOM object to the remote NewsSubscribe code. This code adds the name and email address into a database. At this point, the server can send **any** update back to the device it wants. It could download additional parts of the application if it wanted to, reconfigure the user interface in any manner, etc. However, in this case, we just want to change the current screen attribute from screen1 (the submit screen) to screen2 (the thank you screen.) The server will create a result tree that performs this modification. The entire network message will look like this. The boldface item shows the desired change.

```
<sqml>
45 <network-control seq=3 type="http" action="update"
    terminalid="A23BE3FEC23223A"
    sessionid="B44023957FA992E4EE"
    application="news" />

    <application id="news">
50      <gui id="gui" currentScreen="screen2"/>
    </application>
```

</sqml>

The client receives this clipped tree and folds it back into the superstructure. The screen changes from screen1 to screen2 and the user is shown a “thank you” message.

- 5 It will be understood that the implementer is free to develop his or her own methods of securely transmitting clipped tree fragments back and forth across the network. In other network update models discussed herein, instead of exchanging fragments, the two communications partners could exchange entire objects stored within the superstructure. In this case, sending an object from one to another
- 10 instantiates that object in the receiving superstructure, replacing any previous definition of that object.

VIII: Security

- The following sections survey various technology aspects of SimpleOS including
- 15 security, stylesheets, and embedded systems.

Security Features

Runtime Security

- Several fundamental design aspects of the SimpleOS container ensure that it is inherently resistant to almost all of the runtime security issues that plague modern
- 20 operating systems such as Windows. There are two prongs to this advantageous design – the enclosed functionality and the presence of an intermediate form. These two advantages give SimpleOS a level of security similar to that of other enclosed languages like pure HTML, Java applets and Flash.

Enclosed Environment

- 25 Enclosed execution models, such as SimpleOS, HTML, Flash, and Java applets, are inherently secure because they allow the programmer to do almost nothing outside of a well-defined container. In HTML without JavaScript, the only ability is to present a page and cause images to be downloaded. Java applets cannot access the local file system,

interact with other processes or systems, issue spurious network connections or open unwanted windows. Although there are exceptions in the case of Flash, even with this technology it is relatively true that the only capabilities allowed are restricted network communication and the guarded ability to make animations and multimedia presentations.

SimpleOS follows in the tradition of programming languages that discard general features in return for greater security. Within a SimpleOS application, the only manipulations possible are those within the application's superstructure. This is confined to managing a user interface with pre-defined controls, playing sounds and animations and performing certain authorized telephony and network communication operations. Tree manipulation is the extent of the ability of a SimpleOS application. An analogy would be a web page that contains nothing but HTML and the ability to only change the HTML. In this scenario, it is virtually impossible to perform a bad operation. The worse that could happen in a situation like this is that the application runs away in an infinite loop. The worst that can happen in a situation like this is that the application runs away in an infinite loop.

Unlike a general-purpose language such as Visual Basic, C, C++ or J2SE, there simply isn't the means to express inside the language any operation outside of its specification. If SimpleOS does not explicitly provide for an operation to occur, it is impossible for the application to find a way to do it since the application relies entirely on the superstructure and SimpleOS to perform all of its work. This is not true in C or C++, which can run code directly on the processor and access thousands of API toolkits that can perform various operations. Because languages like C/C++ and Visual Basic operate in environments that are meant to give the programmer general-purpose functionality, it is very easy to find holes in the tens of thousands of features inside the environment. SimpleOS, not being a general-purpose environment, does not suffer from the proliferation of toolkits and their concomitant security risks. Incidentally, SimpleOS does allow Turing-equivalent general computation – it simply restricts access to the environment from the components capable of computation.

C and C++ also suffer from the difficulty of being processor native, which means that they have full access to main memory, registers and other aspects of a system close to the hardware. Operating systems must specifically cordon off C or C++ application from unwanted activity, and many security holes appear because the cordon is not as tight as it should be. These languages must be confined by a negative security model that relies on restraining activities. Interpreted environments such as SimpleOS rely on a positive security model where only the operations that have been considered by the designer are allowed inside the language.

C/C++ Application toolkits present a high security risk because there are so many of them and they frequently pass data from caller to caller in a low-level binary format. SimpleOS does not provide any API toolkits at all, and does not suffer from this risk in the same way as a binary runtime environment.

SimpleOS also provides a security advantage in that it does not allow any low-level processor operations to occur. The only direct work possible from within a SimpleOS application is the rearrangement of trees.

Intermediate Form

The second major security aspect of the SimpleOS design lies within its intermediate application data structures. In SimpleOS, there is no direct correlation between program code and anything from the host environment. This provides a level of security that would normally be found only in static languages such as non-JavaScript enabled HTML. SimpleOS is unique among almost any other dynamic language in the strength of this separation.

Even in a language as well protected as Java, emulated virtual machine instructions are eventually translated to corresponding native operations. This translation exposes a potential risk that bad data generated within the container might be copied without protection into systems outside of the container. This can expose weakness in the API toolkits of the host environment to malicious code within a container. This type of problem is difficult to contain since the translation typically occurs in a wide variety of places within the emulator.

The key to SimpleOS's data separation between program code and guarded API calls lies in the structure of the SimpleOS execution model. Within SimpleOS, an application's only possible operations are tree transformations. The immediate result of any executed code within SimpleOS is never more than a modification to the application's superstructure. Only after the code has completed does SimpleOS engage in a synchronization task, where changes made are carefully extracted and API calls issued the host operating system. During this pass, SimpleOS has the complete freedom to examine and restrict the application before the host operating system is ever invoked to perform an update. This extra pass provides a unique security advantage over other languages.

Once again, a SimpleOS application only has control over an intermediate form, and cannot touch the underlying interface that allows communication with the user and other devices. SimpleOS always makes the actual manipulations within the host operating system on behalf of the SimpleOS application.

The intermediate form found between the application's code and the host operating system is entirely based on an XML Document Object Model. There is very little possibility for rogue data to be added to this model since all data types are explicitly marked. A united set of programming interfaces governs the extrication of data during the rendering process. Because all arguments are extracted from a single common model, it is easy to provide a central clearinghouse within the implementation to check data before it is used in the platform-native sections of the SimpleOS code. HTML also benefits from this strong separation, and that is why 100% pure HTML rarely has been known to provide a security fault. When such faults occur, they are easy to identify and repair. Just like SimpleOS, HTML never produces an immediate list of instructions in a host environment when it is loaded. Instead, an intermediate form is produced first. This intermediate form is examined and used by a rendering engine to produce a final form. Any implementation of software is subject to development defects that could expose a security weakness. An effective approach for dealing with security issues is to address them by design, and the SimpleOS design methodology uses the very best defenses against programmatic error. By providing functionality

containment and an intermediate form, SimpleOS is one of the most secure programming environments.

Incorporation of Additional Security Features

The present invention also allows for a number of additional features to implement security, which are useful, among other applications, in addressing the security concerns of an operating system run on wireless devices. There are several classes of security threats and means to address them, as shown in Table 53.

Threat Type	Description	Security Approach
Interception of sensitive data	Intruder obtains sensitive information such as credit cards, personal information, accounts numbers, etc.	If the communication medium used does not support encryption, SimpleOS will use a security library to provide RSA encryption.
Man-in-the-middle Attack	Intruder causes a client's information be redirected through a compromised system that attempts to modify or insert unwanted code over the network.	Messages to the device must be signed with an MD5 digest that is verified before the content is used.
Buffer overflow attempts to overwrite sections of SimpleOS code	Intruder runs code that uses string arrays that exceed an internal buffer, causing code sections to be overwritten.	Not possible (or at least, extremely unlikely) for Java based architectures. In non-Java architecture, the only buffer overruns possible are in XML data, so a central place to verify buffer checking can be found in the code. This will reduce the risk of buffer-overflows.
Fatal programmatic errors	Intruder runs code that attempts to divide by zero or perform some other functional impossibility.	Tight checking of the expression evaluation code. Using existing, well-reviewed code for the expression language (for instance, longstanding commercial implementations of XPATH) will reduce risk considerably.
Infinite Loops, Stack Overflows	Intruder runs code with looping constructs that do not terminate or runs an unbounded recursive function.	Limited ability within expression language to create unbounded loops. Because all data structures are part of an XML parse tree, an excessive number of repetitions of data structures will cause application errors within the XML implementation.
Unwanted device interaction	Intruder runs devices that dial unwanted phone numbers, erase address books, etc.	Certificate model will prevent access to anything outside SimpleOS without an authorized certificate.
Unexpected API interaction	Intruder writes code to manipulate the API in such a way that SimpleOS crashes or causes unwanted behavior	Extremely unlikely because the only action possible by the programmer is to modify a tree. (See discussion of intermediate form below.)
Unwanted native code plug-ins	Intruder writes a native code plug in that performs unwanted behavior.	A certificate system will protect any dynamically linked native code plug-ins.
Viruses	Intruder writes self-replicating code.	A certificate will be needed to access other applications besides the one provisioned on the device. Consequently, the ability for one program to infect another is impossible unless a rogue certificate is acquired. In addition, since SimpleOS will always prompt the user to allow device-to-device communication, the spread from one device to another is unlikely.
Rogue Certificate	Intruder illegally acquires a code signing certificate and uses it to distribute code that performs unwanted behavior	Eventually a certificate revocation system can be implemented within SimpleOS to provide an additional level of verification for a certificate
Trojan Horses	Intruder writes an application that appears to perform one task, but is secretly performing another.	Applications are limited in the activities they can perform outside of the SimpleOS container, but if a mobile operator wants to be particularly careful, they can prevent SimpleOS from installing any non-authorized programs.
"Spyware"	An application sends back private data to a central server without the user's knowledge.	If desired, network communication can be configured to require a special authorization certificate.

Table 53 - Summary of Security Threats and Responses

Hash Signing

In one practice of the invention, to prevent security attacks where a rogue server provides unwanted data to an application, all network communication can be secured with a cryptographic MD5 hash. This hash is based on a shared private key that is
5 exchanged when the application is first provisioned. All messages exchanged between the device and the server must be signed by this hash, and any message exchanged between applications must contain the application's hash.

MD5 is known as a cryptographic digest. When the algorithm is run over a message, the function produces a short signature. Even the smallest difference
10 between two nearly identical messages will cause a vast difference in the value of the digest. The difference is impossible to predict in advance without resorting to lengthy combinatorial code-breaking, so the hash signature does an excellent job of assuring that a message contains its original content.

At application provisioning time, the server and the client set up a private key
15 shared between them consisting of several hundred bytes of random numbers. Before a message is exchanged, the sender computes the MD5 hash of the concatenation of this key and the message. It sends the resulting hash code and the message together. The receiver adds the private key to the received message and checks to ensure that the code is the same. In this manner, the MD5 signature provides a reasonable guarantee
20 that the message has not been tampered with and that it came from a source that knows the private key.

Encryption

Most transmission protocols used by wireless carriers maintain some sort of pre-existing encryption mechanism. Because SimpleOS will use the communication means available
25 to it from the host operating system, this will cover protecting the data from unwanted interception.

In cases where the host does not provide security, Simple Quick GO has investigated several very lightweight implementations of SSL (Secure Socket's Layer)

available in C and Java. Such an implementation will be used to secure data transmission.

Certificates

5 In one practice of the invention, depending on the needs of the mobile operator who deploys SimpleOS, applications can be restricted in their operation unless they were provisioned with a cryptographic certificate. In such a scenario, each application provisioning server would be required to provide a class-3 digital certificate, signed by Simple Quick GO in order to secure and encrypt the transfer of an application to a device. Furthermore, the application may be limited to certain functions based on the
10 type of certificate they are provisioned with.

Functions that can be optionally limited for a particular install base include:

- Network access
- Access to phone and telephony features such as the address book and initiating
15 a phone call
- Access to other applications deployed in SimpleOS
- Ability to use local area network communication (BlueTooth, 802.11B)
- Ability to load and use dynamic code
- Ability to install and run an application
- 20 • The ability to do any of the previous without explicit user confirmation

Certificate technology has been used successfully by Java applets and Microsoft ActiveX technologies to limit the ability of unwanted code to run on a device. A detailed description of digital certificate technology is beyond the scope of this document, but a
25 brief description follows. A central authority responsible for an installed based of SimpleOS deployments maintains a certificate server and an issuing office. This central authority (CA) might be a carrier, mobile operator, device manufacturer or Simple Quick GO itself. Developers approach the CA asking for a particular certificate (such as a certificate to perform telephony.) The CA will issue this certificate to the developing

organization, which uses the certificate to cryptographically sign each application. This signature is recognized by the SimpleOS container and verified before performing a restricted activity.

To promote the widespread expansion of the SimpleOS container, an overly restrictive policy may not be recommended. Only the most dangerous operations such as installing an application or adding dynamic code might require certificates. It is expected that mildly risky activities would require a simple confirmation from the user before they are performed. In the case of an action requiring a certificate, user confirmation will not be required.

Certificates are used to maintain a chain of responsibility. If code that has a valid certificate is found to contain a virus or other dangerous code, the owner of that certificate can be traced. Since the certificates are hard to forge, it allows an operator to have cryptographic certainty that only authorized developers are performing dangerous operations on a device.

Applications can sometimes be provisioned in a peer-to-peer arrangement instead of from the server. In this case, the certificate is transferred along with the application inside the resource stream that provisions the application.

IX. Stylesheets

While stylesheets are not necessary to a basic superstructure-based application environment, their use is of great utility to producing a highly functional implementation of the invention. Accordingly, one embodiment of SimpleOS employs a style-sheet mechanism for the definition of colors, fonts and other presentation attributes of user interface elements. This allows the presentation of the GUI to be separate from the description of the GUI components. This same approach has been successfully used in other areas of web technology, notably HTML.

A stylesheet groups sets of formatting properties such as fonts, colors and layout into a style classes. Each element in the GUI is associated with a style class, which provides that element with its presentation attributes. (A default style class is associated with any GUI element that does not have an explicit class.)

The example of Table 54 shows a very basic SQML application that uses a default stylesheet to establish parameters for a set of frame buttons. Each frame button will inherit from the class since it is marked "default".

```

5  <?xml version="1.0" encoding="UTF-8"?>
   <!DOCTYPE sqml SYSTEM "sqml.dtd">

   <sqml appid="com.sqgo.simpleos.application.launcher">

10     <sq-style-sheet>
        <class default="true">
            <palette>
                <colordef name="blue" rgb="#4444ff"/>
            </palette>
15     <bgcolor color="blue"/>
        </class>
    </sq-style-sheet>

    . . .

20    <sq-card id="card01">
        <top label="TOP" id="top"/>
        <left label="LEFT" id="left"/>
        <right label="RIGHT" id="right"/>
25    <bottom label="BOTTOM" id="bottom"/>
        <frame id="frame01">
            </frame>
    </sq-card>

30    . . .

</sqml>

```

Table 54 - Simple Stylesheet Example

SimpleOS stylesheet classes exist in an inheritance hierarchy, which allows one class to extend another. This allows a developer to quickly modify the look and feel of an application. For instance, the developer could define a base stylesheet for a set of applications, and create sub-classes of that stylesheet for each particular application.

40 This would allow an attribute, say the frame color, for all of the applications to be automatically the same. If the developer wishes to change the color of a particular application, he can modify the sub-class of that application's stylesheet and still maintain other properties mapped from a master sheet. The same pattern applies to screens within an application or controls on a screen.

The mapping of each component to its ultimate style is actually performed at provisioning time. The provisioning server creates a master map of all of the styles in use by an application and emits a set of condensed, non-inheriting style classes as part of the application. This expansion conserves the amount of processing that must be done by the deployed application on a recurring basis. This optimization is only effective if the application's code remains constant through the life of the application. Since our model does allow additional application code to be sent during the evolution of the process, the SimpleOS container will also contain logic to update stylesheets that use inheritance.

Table 55 shows an example of a more complex stylesheet.

```
<sql:stylesheet>
<sql:class id="base">
  <sql:palette>
    <sql:colordef name="white" rgb="#ffff00"/>
    <sql:colordef name="black" rgb="#ffff00"/>
    <sql:colordef name="red" rgb="#cc0000"/>
    <sql:colordef name="blue" rgb="#4444ff"/>
    <sql:colordef name="yellow" rgb="#ffff00"/>
  </sql:palette>
</sql:class>
<sql:class super="base" default="true">
  <sql:bgcolor color="white"/>
  <sql:font typeface="Helvetica" typesize="12" typecolor="black"/>
</sql:class>
<sql:class super="base" id="border">
  <sql:bgcolor color="blue"/>
  <sql:font typeface="Helvetica" typesize="16" typecolor="white" typestyle="bold" />
</sql:class>
</sql:stylesheet>
```

Table 55 - A Complex Stylesheet

As noted above, stylesheets provide the mechanism for managing screen layouts on different devices. One of the properties that can be contained in a style is an area indicator that will give the exact width, height and location information for a particular GUI element.

X. Additional Features

Embedded Applications

SimpleOS is designed to allow a fully enabled implementation of the container to communicate with embedded devices. For this purpose, embedded devices may include

TV set top boxes, vending machines, home appliances, car systems and other portal devices. There are a variety of adaptations to the system that apply to embedded communication that are discussed in the sections below.

Simple Embedded Protocols

- 5 Some embedded situations call for sending and receiving well-defined codes between the device and a remote embedded system. These codes may be transmitted using infrared, radio frequencies or a serial connection. Typically this type of communication uses a fixed vocabulary of signals. SimpleOS can accommodate this type of interaction in several ways. One option is to provide a toolkit accessible using a limited native
- 10 interface from within SQScript. In this situation, Simple Quick GO, or some other author, would provide a SQLML interface to the communication pathway required. Incoming messages to the SimpleOS-enabled device would be directed to the message queue of the application and delivered via a message descriptor. Outgoing messages would be posted by filling in a tree-object as a request and posting it to the operating
- 15 system using a primitive within SQScript. If the pathway ends up being very common to a class of users, Simple Quick GO may even add support for it directly into the operating system.

Code sets to control a remote device might include:

- 20
- A vocabulary of IR codes to control a television
 - A command set for an internal automotive control network
 - Radio frequency mappings for a remote-controlled car
 - Diagnostic codes available from the OBC of a car engine

Network Protocols

- 25 In more advanced embedded applications, the two devices in question (the SimpleOS-enabled device and the remote device) form a network based on some existing protocol. These scenarios might include PPP over IR, TCP/IP over WiFi, or TCP/IP over Bluetooth. In the case that the communication path is packet-based, a native plug-in

may be created just as described in the previous section. In this case, the data packets that are sent and received are channeled to an application using code added into the operating system.

In the case where a streaming protocol is available, the approach is much less complex. In this case, SimpleOS can communicate directly with the device using HTTP or HTTPS. This has proven to be very effective in embedded devices since a TCP/IP stack and web server can usually be implemented in less than 4K of program code. In this situation, messages can be passed to the web server using the same network communication model used for server-side communication. If the embedded device can stream enough data, it is even possible that it can provision a new application to SimpleOS. In this case, the application would have to have been completely pre-provisioned for the specific device type, and the end application bundle signed and certified. The application would simply appear to the embedded system as a stream that can be shipped over to SimpleOS.

15 **Vending Machine Example**

Using a simple wireless network and a streaming protocol could allow for some very powerful results. Imagine a vending machine that allows a cell phone to connect to it to receive a menu and verify payment. The exchange may operate as follows:

- 20 1) The user approaches the vending machine and enters the wireless network managed by the vending machine's computer
- 2) During the polling of the new network, the launcher application running inside SimpleOS discovers that there is a provisioning source available on the local network. A message is displayed to the user: "You are near a vending machine that can be controlled from your cell phone. Would you like to connect?"
- 25 3) SimpleOS sends a message to the vending machine with its device characteristics. The vending machine already has a set of pre-provisioned, signed applications available to control it, and it checks to see if the cell phone's device type is listed among the pre-provisioned images it has.
- 30 a. If the vending machine does have the correct image, it streams it over to the cell-phone and provisions the controller application
- b. If the vending machine does not have the correct image, it sends a URL to a provisioning server to the launcher application. The launcher application can then download the correct software.

- 4) The launcher application, having provisioned a new application, asks the user: "You have just provisioned a new application. Would you like to run it now?"
- 5) The application loads, and presents its user interface.
- 6) The user chooses an action such as "vend" inside the application
- 5 7) The target of the message descriptor includes a variable that causes the message to be marshaled as SQML and deployed to the vending machine over the wireless network.
- 8) A simple HTTP engine on the vending machine parses apart the message and performs the action.

10 **SimpleOS to SimpleOS**

Two SimpleOS devices can easily connect to each other and accomplish a variety of communication models. The simplest is when one SimpleOS instance connects to another for peer-to-peer provisioning. If the two devices have different device type signatures, this may require the receiving instance to connect to a remote location to
15 receive a device-customized version of the application. Otherwise, the provisioning will work securely because the certificate has been embedded in the resource stream, guaranteeing to the receiver that the application has not been tampered with.

Two SimpleOS instances capable of exchanging messages can place items in each other's queues. This is made possible by the fact that all application event
20 messages are expressible as SQML. If an application running on one device knows the address of another device, it can formulate a SQML message using the standard remote messaging model and deliver that message into the queue of the receiving application. This model allows for several compelling uses of the technology. For instance, a TV set might itself run an instance of SimpleOS that has linked into it commands to control the
25 low level systems inside the TV. Any other SimpleOS device that can connect to the SimpleOS system inside the TV can post messages to the TV without any additional development work. The two operating systems speak the same language; so one of them can serve as the adapter to a particular piece of hardware on behalf of another device.

Benefits and Features

The SimpleOS design presents a radically new way of organizing a computer program, coupled with an innovative graphical user interface design. Applications are freed from the need to understand how to manipulate their outside environment. Instead, they
5 update organized descriptions of their current state. The operating system does the work to determine the specifics of how to implement the request. By doing the heavy lifting for user interface and application logic, and creating a single development language, the SimpleOS model offers easier development of widely deployable applications. Also, the powerful network features within SimpleOS greatly reduce the
10 amount of work to make an application network aware.

All of these factors become particularly important when considering the limited abilities of most mobile operating systems such as J2ME and Brew, and the fact that code written in one environment is almost useless on the other.

Streaming

Given the superstructure-based application environment of the invention, a system could be readily implemented (see FIG. 2) that would "stream" updates to the superstructure as a broadcast to multiple SimpleOS devices at once. Such a system could utilize either (1) a means of ensuring that each device was in a consistent, known
20 state at the time of the transmission and that the transmission remains whole and complete; or (2) transmitting complete segments (each all or nothing) of application update.

As shown in FIG. 2, a method in accord with this aspect of the invention includes broadcasting update data objects to devices, while ensuring coherency (202); receiving,
25 in the devices, the update data object(s) (204) and updating in accordance with the update data object (206). Note that the updates could include updates to applications, data, or both.

Consider, for example, the example of a ballpark filled with baseball fans, many seeking dynamically updated information about scores from other games, or
30 pitcher/hitter statistics (how has this pitcher fared against the next batter in the past?).

Using the streaming update feature, updates could be broadcast to thousands of SimpleOS-running cellphones or other handheld devices at once, using either of the two noted techniques (either a synchronization/coherency-insuring technique, or the transmissions of all-or-nothing segments of application update, so that users could receive real-time updates of an application displaying hitter/pitcher statistics and other information. If particular fans do not have the necessary application on their handheld device, they could be queried as to whether they'd like to download the application (instantiation as a form of "update") and the methods described herein could be used to instantiate the application, and subsequent updates, into the handheld device.

These transmissions could use a data broadcasting method that could minimize the need for individual point-to-point data transfers between the update server and each client. The invention thus enables update streaming without the requirement of sending a multiplicity of different messages on a per-client basis for each update.

In each case, the update would be incorporated into the superstructure using the superstructure modification/update techniques described in the earlier sections of this document. Because of the superstructure-based application environment of the invention, the updates could be efficiently transmitted and processed across a wide range of different handheld devices within the stadium, regardless of the native operating system of the various processors involved.

Streaming in this manner provides a number of advantages, including:

- (1) less per-client state is required on the information server;
- (2) the technique can exploit the network efficiencies of known or new multicast/broadcast technologies, and reduce network traffic;
- (3) it reduces the need for retransmission of failed packets (the next broadcast would overwrite them);
- (4) it allows device transmitters to remain in "receive only" mode instead of having to send back confirmation; and
- (5) there is no need to maintain records of precisely which devices are still receiving information, so devices are more free to leave and re-join the network without interruption to the communication protocol or the user experience.

Equalization

Mobile applications are typically divided into three areas:

- 5 • **User Interface Code:** The look, feel and presentation of the application
- **Application Logic:** The specific behaviors or features of the application
- **Communication Logic:** The parts of the application that communicate to the outside world to receive or send information

Depending on the platform, typical application development can focus greater than
10 80% of the attention on user interface logic and communication. Many mobile operating environments are particularly constricted, and require highly specialized programmers to work very carefully to create easy-to-use interfaces and manage communication. This is unfortunate since most of the core aspect of the application is not present in either of these areas. Worse, the work done to solve these problems on one platform
15 can almost never be reused when porting to a different platform. Together, these issues increase testing and development time (and therefore, development cost). Given that the market for wireless applications is still in its infancy, this acts as a huge disincentive for developers.

The left-hand side of FIG. 20 shows the traditional application model (2002) that
20 would commonly be applied to a mobile platform such as Qualcomm's Brew (Binary Runtime Environment for Wireless.). Brew supports a very small subset of features for communication, data management and graphical presentation. These features are presented to the application programmer through a programming interface that requires procedure calls to set up and maintain hardware resources. The application
25 developer typically must perform additional work to tailor the raw capabilities granted by the host environment into the specific look and feel, communications model and data representation required for the application. On top of this initial effort, the programmer adds the specific logic or functionality that the application requires. Most existing platforms for application development generally follow this model.

30 The work done to shape an application's services on a particular software environment typically cannot be re-used between environments due to dissimilarities

between different platforms. While the top-level application logic can occasionally be re-used between toolsets, this only applies to platforms based on the same language such as C or Java.

In contrast, a SimpleOS product implemented in accordance with the invention
5 (2004) sits on top of each target operating system providing the user interface, communication and data organization features. SimpleOS, controlled entirely by an application's XML structure, insulates the developer from the specific details of the platform.

This model extends well-beyond the "write-once, run anywhere model of Java.
10 Because programs can be described in a lightweight form of XML that can run on a variety of hosts, SimpleOS makes it easy to deploy applications anywhere.

Networking Architectures

SimpleOS provides the ability for server-side processing to occur just as easily as it can on the device. The cornerstone of this feature is that the networking protocol can ship
15 over arbitrary sections of the application's superstructure over the network by converting it into SQML. In essence, SimpleOS is sending application snippets instead of documents or data when it communicates with the server.

This networking paradigm is the natural evolution of user interface-to-server architectures. Originally, mainframe-computing systems relied on a text-based user
20 interface and communicated through specific screen elements. Client/Server architectures supported greater user interface logic local to the device, but were still constrained by communication and desktop-specific programming rules. The web services model uses only one client browser, but all application logic must reside on the remote server. SimpleOS provides for presentation and application logic to transparently
25 reside on either the client or server, providing a leap forward in application design. See 2108 of FIG. 21, contrasted with 2102, 2104, and 2106.

Table 56 illustrates the differences in communication and data paradigms.

Architecture	Mainframe	Client/Server	WWW (including WAP)	Simple Quick Go
Central Communication Paradigm	Text Terminals	Row sets and Actions	Documents	Applications
Transport Atoms	ASCII	SQL	HTML	SQLML
Reliance on communication	Extreme – without network, client is dead	High – Client is usually useless without server	High-Medium – Browser can view documents offline, but interactivity is impossible	Low-Medium – Applications have the option to spread interactivity between client and server, and applications can be standalone
Device Specific Programming	None. Programmers just display ASCII	Clients must be written to native platforms	HTML/WML is general, but functional versatility is limited and tied to the server.	None. SQLML offers maximum functional versatility and client-device independence.
Location of application logic	Server	Server	Server	Client and Server in any proportion
Location of data	Server	Server	Server	Client or Server as needed
Description of user interface handling	Server	Server	Server	Client or Server as needed

Table 56 - Comparison of Remote Computing Paradigms

5

Yet another powerful advantage of the SimpleOS network model is that an application's code can be updated on the fly if so desired by the carrier or the developer.

10 **Platform Comparisons**

The following sections compare SimpleOS to several existing technologies to help illustrate the differences in the technology.

SimpleOS vs. Java

15 SimpleOS and Java share many common goals. Both operating systems are designed to run as virtual containers over host operating systems like Windows or Linux. Both operating systems allow for write-once, run-anywhere. Both provide a consistent platform-independent abstraction between their running applications and the native operating system. Both are also highly network-aware, object-oriented and have support for XML.

The major difference between SimpleOS and Java is that Java does not support the novel execution model of SimpleOS described herein, which frees the programmer from concerns about how a particular description of an application should be implemented, and allows a purely descriptive development model. In SimpleOS, the application need only describe changes. In contrast, within Java, the application must also perform the changes. XML support within SimpleOS is automatic – any running application can communicate in XML as easily as it communicates with itself. In contrast, Java requires marshalling and un-marshalling of XML data into private data structures. SimpleOS supports an intermediate form of application state that provides an additional level of security beyond the insulation provided by the JavaVM. In addition, SimpleOS's network model allows for a seamless sharing of applications and data between devices, because superstructure may be accessed and transferred. In a Java environment, this same effect would require explicit transfers of data and the programmer would need to identify when those transfers would be required, and write code specifically to handle them. It is expected that SimpleOS code sizes would much smaller, because the application need not be concerned with support for a graphical user interface library. Also, SimpleOS is designed to embed the entirety of the state assumptions between the application and the environment into the superstructure. The same is not true for Java, which stores extensive private data within its class libraries, and even within the device-native portion of the virtual machine. A frozen Java program can not consistently be restarted outside of its Java container. The introduction sections above, which describe other properties of the invention, discuss other differences between opcode-based execution environments and SimpleOS.

SimpleOS vs. HTML

It may initially appear that HTML and SimpleOS have nothing in common, but the truth is that they share common roots. Both of them rely on a complex data structure that is entirely descriptive; SimpleOS has the superstructure, and HTML has the DOM. However, HTML describes documents, where as SimpleOS's SQML describes

applications. A SimpleOS application can be completely client-side and does need to communicate with a remote server in order to be interactive.

At a technology level, both HTML and SimpleOS rely on an intermediate form to gain security, device-independence and stability. However, the HTML intermediate form is relatively fixed once it is downloaded into a web browser, whereas the SimpleOS intermediate form (the superstructure) may change freely over time and is designed to allow such changes to occur as a natural part of its execution. The entire operating "contract" and shared state assumptions between a SBAE engine and its application code is explicitly expressed within the grammar and content of the superstructure.

SimpleOS does not consider the current running application to simply be an XML grammar changing over time after it has been seeded onto a browser, but instead, a rich, dynamic collection of state variables, structured data and program code that can easily be marshaled back into a self-consistent package and transferred to another device. Values within the superstructure may refer directly to expressions that evaluate within the context of the data being read. Additionally, the update mechanism used by SimpleOS works by directly delivering events into the superstructure itself. The superstructure is then free to modify itself in any way possible, as long as all modifications are captured within the superstructure itself. Finally, the design of SimpleOS allows for an application's events and data to be handled seamlessly over a network, in part because the activation of event handlers and the data required to process them has been generalized. These differences, and others, arise as a result of the inherently different design foundation between HTML/DHTML, (optimized for documents), and SimpleOS (optimized for applications).

SimpleOS vs. XSLT

SimpleOS and XSLT share some characteristics in that they both use template languages to effect change within trees. SimpleOS may even use XSLT in some of its early implementations because so much work has already been accomplished within XSLT in the area of tree manipulation. However, XSLT is almost entirely about

producing final form grammars for incoming data. Typically an XSLT processor receives data in one form of XML, and then transforms that data into a different XML grammar. In SimpleOS, the transformation always occurs within the same grammar (SQML.) The output of the grammar does not leave the system, but instead it gets reapplied to the original document. In this respect, SimpleOS's transformation model can be a closed loop where as XSLT is almost always an open loop.

It is to be understood that the invention described herein is amenable to a wide range of variations, modifications and alternative constructions and implementations. It should therefore also be understood that there is no intention to limit the invention to the specific implementations described above. On the contrary, the invention is intended to cover all modifications, alternative implementations and equivalents falling within the scope and spirit of the invention as defined by the claims appended below.

Glossary of Terms Used in This Document

API: Application programmer interface. The means for a traditional application to change its outside environment and communicates with the operating system.

Class: A template that can be used to instantiate objects.

Clipped Tree: A form of a result tree that minimally describes a set of changes to an existing tree.

Data Member: A variable declared directly underneath an object or class

Data Tree: The "data" section of an object, which can be an arbitrary SQML tree.

Dynamic Section: A portion of the superstructure that has shared meaning between the operating system and the application. When external events occur, the operating system updates dynamic sections inside the superstructure and notifies the application. When the application wishes to change its environment, it updates the dynamic sections of the superstructure and the change is observed by the operating system.

Element Identifier: Syntax for referring to an individual element within an object

Event Handler: Code that handles an event

Event Target: The location of code that is executed when an event occurs

Host Operating System: The native operating system of a device, such as Windows, Java or Brew. SimpleOS is written as an application or library for a particular host operating system.

Initializer: A method within an object that is executed when the object is instantiated.

5 **Interpolated Expressions:** An expression found in a string that is evaluated and replaced with the value of the expression

Object: An instantiated object containing a data tree, methods and data members.

Result Tree: The output of a transformation-based SimpleOS message handler, which represents all or part of a new tree that must be applied to the superstructure.

10 **SimpleOS Container:** An application that implements SimpleOS and provides a protected area for SimpleOS applications to run.

SimpleOS: A lightweight operating system from Simple Quick GO that runs on a variety of environments such as Windows, Brew, J2ME, and J2SE etc.

15 **SQML:** An XML-based markup language used to describe applications and portions of applications for SimpleOS.

SQScript: A language written in SQML that can be used to express the application logic of a SimpleOS executable. The language is optimized for tree operations against the process's superstructure.

20 **Stylesheet:** A mapping between a display "class" and a set of low-level attributes such as font, color and text size.

Superstructure: A central tree-like data structured maintained by each running SimpleOS process that contains all of the state, code, data and registrations for the entire application.